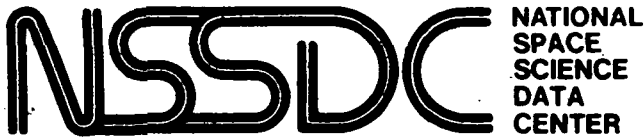


1N-61-TM

32983

054



WORLD DATA CENTER "A" FOR ROCKETS AND SATELLITES

91-06

# Software for Optical Archive and Retrieval

## (SOAR) User's Guide

---

**Version 4.2**

(NASA-TM-105039) SOFTWARE FOR OPTICAL  
ARCHIVE AND RETRIEVAL (SOAR) USER'S GUIDE,  
VERSION 4.2 (NASA) 54 p CSCL 09B

N91-30739

Unclas  
G3/61 0032983



National Aeronautics and  
Space Administration

**Goddard Space Flight Center**

# **Software for Optical Archive and Retrieval**

## **(SOAR) User's Guide**

**Version 4.2**

**Charles Davis  
ST Systems Corporation  
Lanham, Maryland 20706**

**April 1991**

# Contents

1.0	Introduction .....	1
2.0	What Is SOAR? .....	3
3.0	Hardware Requirements .....	5
4.0	Software Description .....	7
5.0	Optical Disk Life Cycle .....	13
6.0	The SOAR Command OCLOSE .....	17
7.0	The SOAR Command OCOPY .....	19
8.0	The SOAR Command ODISMOUNT .....	23
9.0	The SOAR Command OINITIALIZE .....	25
10.0	The SOAR Command OMOUNT .....	31
11.0	The SOAR Command OSECURE .....	35
12.0	The SOAR Command OSTATUS .....	37
13.0	The SOAR Command OTRANSFER .....	39
14.0	Use of the /NULL Qualifier .....	41
15.0	Use of VAX VMS Commands .....	43
16.0	Magnetic Emulation Mode .....	45
17.0	Command File for RCT Formatting .....	47
18.0	Software Installation .....	49
19.0	Hardware Installation .....	51
20.0	Troubleshooting Guide .....	57

The optical disk is an emerging technology. Because it is not a magnetic medium, it offers a number of distinct advantages over the established form of storage, advantages that make it extremely attractive.

They are as follows:

- a. The ability to store much more data within the same space. Each removable 12-inch cartridge, depending on the manufacturer, holds from 1 to 3½ Gbyte of data on each side of the disk. It would take 26 1600-bpi tapes or 3200 floppy disks to do the same job. This capability could solve the ever-increasing storage problem for most companies, and it also makes feasible the transportation of large quantities of data for scientists.
- b. The random access characteristic of the Write Once Read Many (WORM) optical disk. The WORM technology can help in storing large data bases for which random access is a critical requirement. It is also more convenient.
- c. The optical disk has demonstrated a much longer life than that of traditional storage media like magnetic tapes. It is much more reliable under adverse conditions, and external factors like magnetic force, heat, and water have very little effect on it.
- d. The data access rate (read rate: >200 kbyte/sec) is much greater than the rate for magnetic tapes.

The National Space Science Data Center (NSSDC) at Goddard Space Flight Center receives, archives, and disseminates a large volume of space science data, which are accessed repeatedly by scientists. Traditional methods of storage and distribution on magnetic tapes have obvious disadvantages. Much time and money is spent each year to maintain the readability of the tapes. Also, since the scientific community is geographically dispersed, easy exchange of high volume data sets is not always feasible. NSSDC considers the optical disk technology to be a very promising solution to these problems.

7

The write-once characteristic of the optical disk makes it impossible to overwrite or update a block once it has been written to. This may conceivably be seen as an advantage for applications where permanent, unalterable records are required. However, it also makes integrating the optical disk into existing operating systems difficult.

NSSDC has undertaken a project to investigate techniques for integrating optical disk technology into common operating systems. A software interface called Software for Optical Archive and Retrieval (SOAR) was developed to integrate the optical disk into the VAX and the MicroVAX environments. Two possible approaches were considered.

- a. Writing to the disks in a user-designed logical volume format.
- b. Integrating the disks into the VAX VMS system in such a way that data are written to the disks in the standard DEC Files-11 logical volume format.

The second option was implemented using a software interface. By choosing this tactic, the fact that one is using an optical disk becomes completely transparent to the VAX user for read operations. More importantly, once a disk has been completely written to and closed, it becomes independent of any layered software. It can be mounted and read using the standard VMS methods for file access. As long as DEC supports the Files-11 format, which is currently used to write to all the magnetic disks, the data on the disks will be accessible.

The idea behind the design of the SOAR software was that once the scientists fill the optical disks, using SOAR software and appropriate hardware, NSSDC personnel will be able to access and manipulate the data using existing VAX data access methods and applications software. They will be able to use all of the graphics packages and software developed by NSSDC that use standard VMS I/O operations for data access. NSSDC should also be able to send this large amount of data on a few optical disks to its user community in the United States and to overseas users who have access to compatible hardware (specifically the optical disk drives and the host adapters).

The WORM technology does not permit the usual overwriting of directory files that occurs under VMS when files are transferred in Files-11 format. Most of the VMS I/O methods make an implicit assumption that random access disk drives use magnetic recording media. The SOAR software avoids this problem by retaining the directory information in a magnetic working file while writing the actual data to the optical disk until it is filled (or until the user decides that no more data will be written to it). Then it copies the directory information to the reserved lower blocks of the optical disk.



### 3.0

## Hardware Requirements

At least three pieces of hardware are required for an optical disk to be used as a random access storage device. These are

- a. Host computer.
- b. Host adapter.
- c. Optical disk drive.

The host adapter provides a hardware/software connection between the host computer and the optical disk drive.

Most optical disk drives on the market provide an SCSI interface. Almost all of the DEC-supported disk drives provide DEC MSCP protocol. Since SOAR software itself does not provide an independent disk driver (its QSDRIVER is really an extension of a disk driver), it assumes that either the optical disk drive will use a DEC-provided disk driver (like DUDRIVER) or a disk driver will be provided with the hardware.

This fact imposes two requirements for using SOAR software:

1. The optical disk should be mountable
  - a. As a foreign device.
  - b. As a random access disk device.
2. As shown in Figure 1, the host adapter, which interfaces between the VAX and the optical disk drive, must either
  - a. Provide SCSI passthrough between the VAX and the optical disk drive, for example, a UHA-11 controller (TD Systems Inc.) that uses the SCSI bus. In this case a disk driver that supports the SCSI protocol is also needed.
  - b. Convert the DEC MSCP protocol to the SCSI protocol and vice versa (e.g., Emulex UC-14 adapter on Unibus). In this case the optical disk drive uses one of the DEC-provided disk drivers, such as DUDRIVER.

The SOAR software has been demonstrated to support the 12-inch optical disk (1 Gbyte capacity per side) drives made by LMS (formerly OSI), Sony, ATG, and Optimem. Other drives may be supported by SOAR with little or no modification; testing is necessary to validate the software on unfamiliar drives.



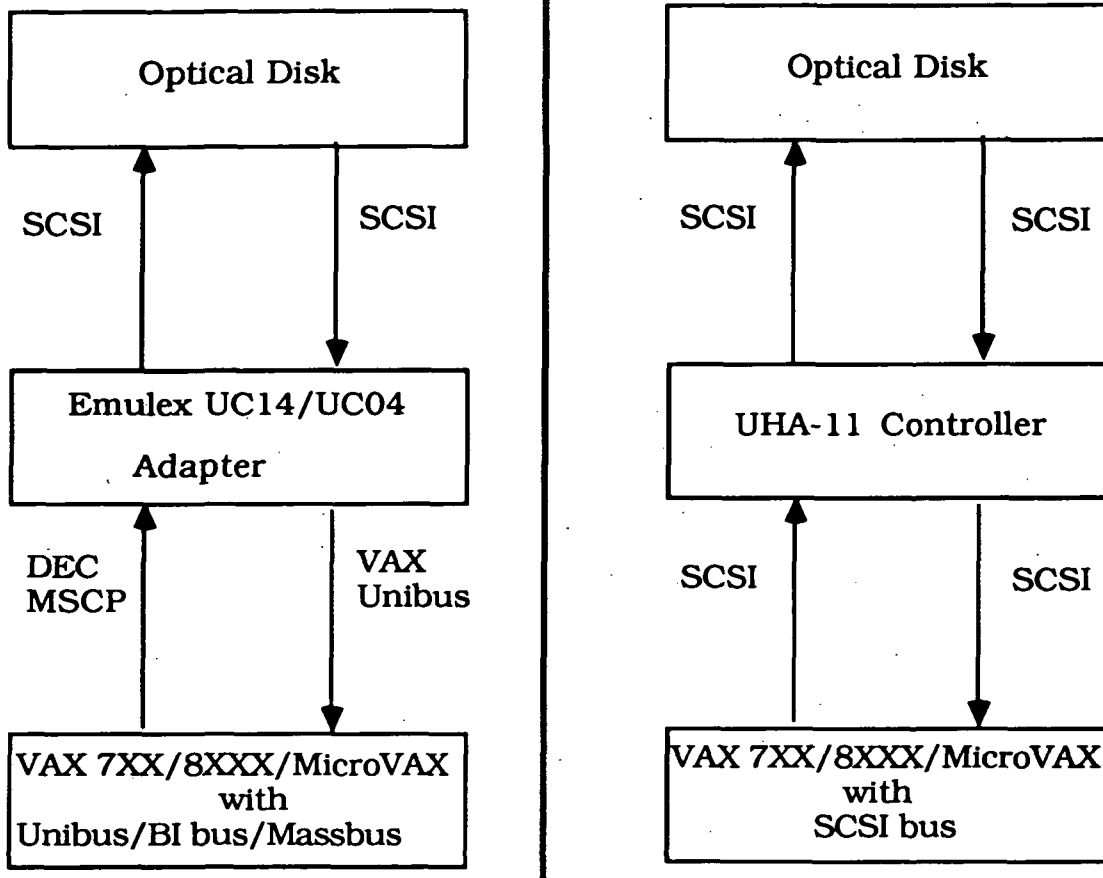


Figure 1. Hardware configuration for VAX/MicroVAX and optical disk interface.

The software consists of a pseudo-device driver and a suite of utility routines. The purpose of the software is to produce a disk that can be accessed via standard VAX VMS Files-11 routines. Once the disk is completely written to, there is no further need for SOAR. During the production of the disk, SOAR's function is to merge a directory file on magnetic disk with the data files on optical disk to produce a whole with the characteristics of a regular VMS logical device. This duality exists until the user has determined that the disk is full. The magnetic work file is then written in the reserved blocks of the optical disk. The disk can then be mounted using the standard VMS utilities, and the SOAR package is no longer necessary. Note that during the dual-media stage, the SOAR logical disk can still be read using the normal file I/O utilities and high-level language calls.

#### **4.1 The Pseudo-Device Driver**

The pseudo-device driver provided with the SOAR software is not an independent traditional disk driver. In fact, it acts as an extension of a disk driver. This means that it performs those functions that the parent device driver cannot handle. Its main function is to connect the magnetic working file containing the directory and the optical disk containing the files into a single logical disk. It does this by redirecting the I/O operations to the optical disk or to the magnetic work file, depending on the type of data block being transferred. For instance, any data block being transferred either to one of the system overhead files (e.g., BITMAP.SYS) or to a directory file is redirected to the magnetic file component of the logical disk. The data blocks themselves, which constitute the user file, are redirected to the optical disk.

Figure 2 explains the function of the QSDRIVER during I/O operations on the pseudo-device.

#### **4.2 SOAR Utility Routines**

SOAR utility routines provide three functionalities:

1. To configure the pseudo-device driver so that the pseudo-device has the same characteristics (disk size, number of cylinders, tracks/cylinder, sectors/track, etc.) as the optical disk and that its extended UCB has the address of the starting block (SBN1) and the size of the magfile.
2. To provide safeguards against overwriting the blocks on the optical disk.
3. To force all I/O operations to start at a physical block boundary on the optical disk; this is necessary because on the optical disk the physical block size is 1024 bytes instead of the 512 bytes used by VMS.

In most cases these routines set up the proper environment and then call the analogous VMS commands. Note that all the I/O operations performed by the SOAR software are on the pseudo-device (QS-device) rather than on the physical device.

Figures 3 and 4 describe each routine and its function during the read and write phases. They also emphasize the dependencies between the routines and the order in which they should be called.

### **4.3 Optical Disk Configuration Under SOAR**

The physical arrangement of the optical disk while using SOAR is as follows:

1. The first magfile-size blocks are a reserved blank area. This area is where the magfile data are written upon OCLOSE.
2. The first block after the magfile area is where the first OCOPIED file goes. Each OCOPIED file is allocated contiguous space, and the files are written one after the other until no space remains on the disk.
3. With the implementation of OSECURE (Version 3.0), the last four blocks of the disk are used to store the magfile attributes.
4. With the implementation of OTRANSFER (Version 4.0), magfile\_size+4 blocks are used to hold any magfiles that have been OTRANSFERRED to the optical platter. These blocks will not be reserved if OTRANSFER has not been executed. OCOPY and OTRANSFER both check to ensure that any existing data are not overwritten.

Figure 5 is a pictorial representation of the physical structure of the optical platter. As can be seen from the diagram, the disk is filled up from the beginning with the magfile reserved area and the OCOPIED files, and down from the end with the OSECURE and any OTRANSFERRED magfiles.

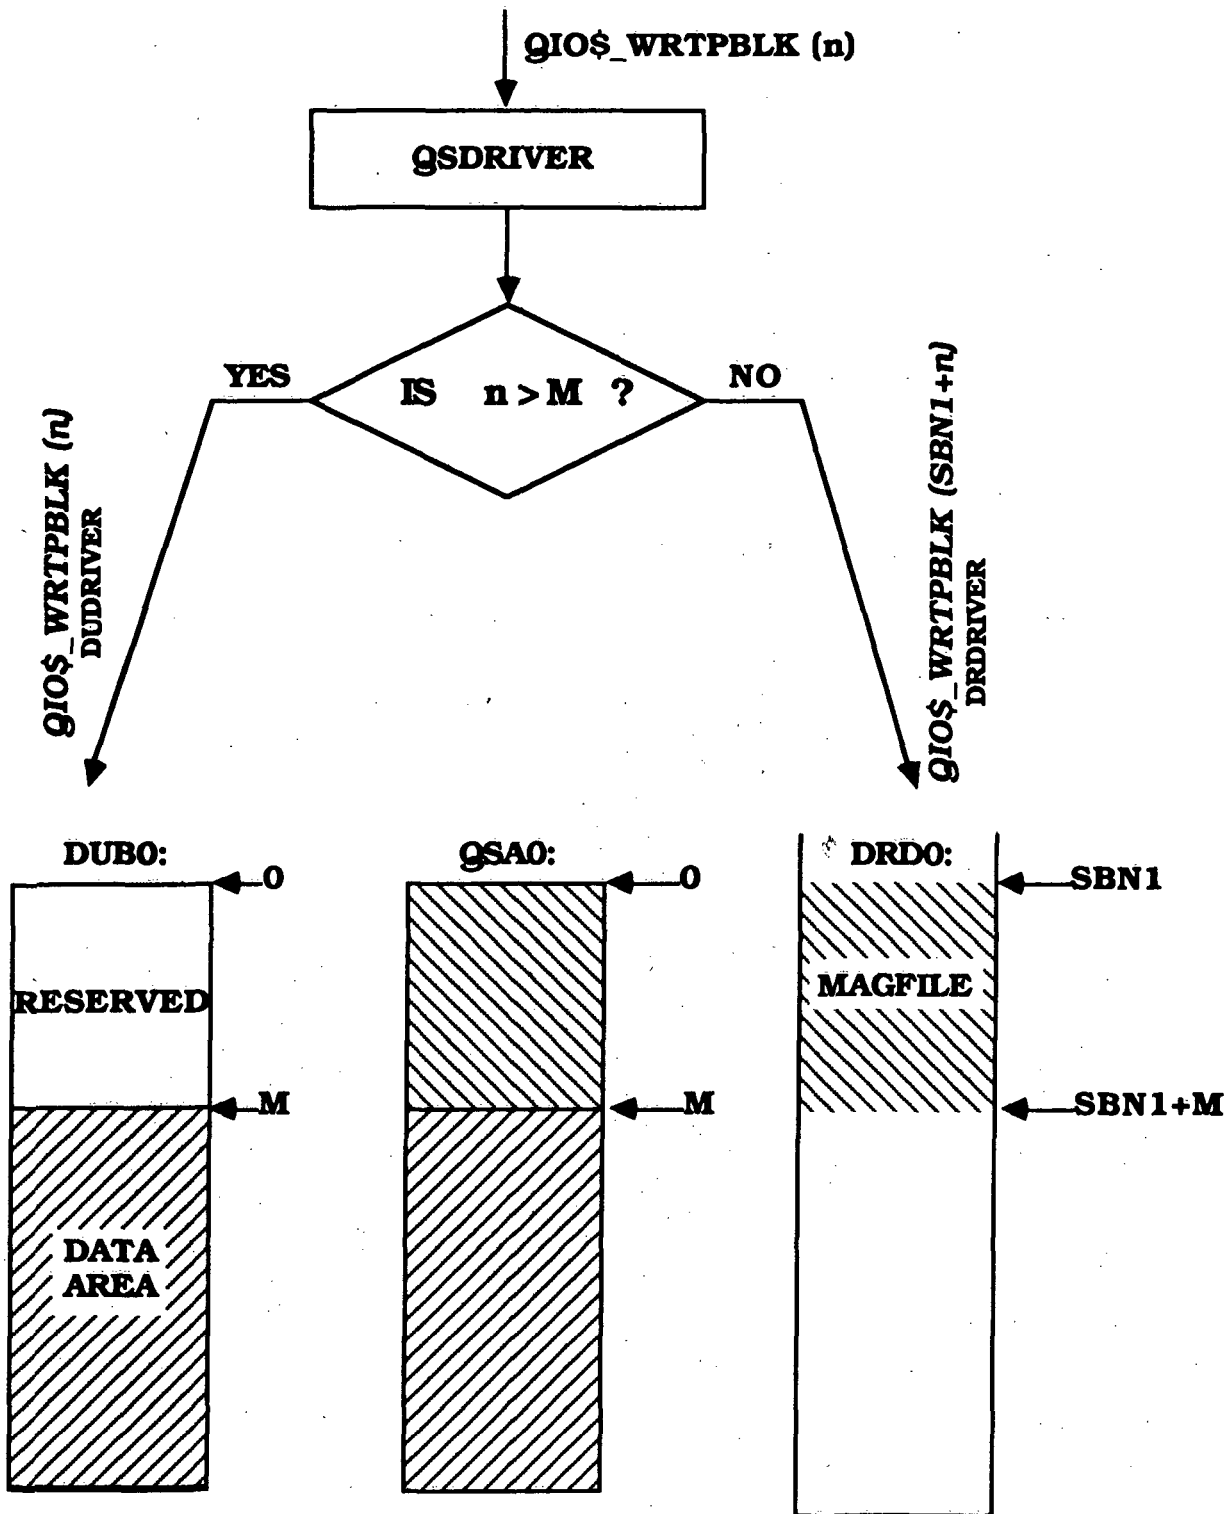


Figure 2. Function of QSDRIVER.

DUBO = optical device      QSAO = pseudo-device  
 DRDO = magnetic device      M = size of magfile in blocks  
 SBN1 = starting block number of magfile on magnetic disk  
 n = block number of block to be written

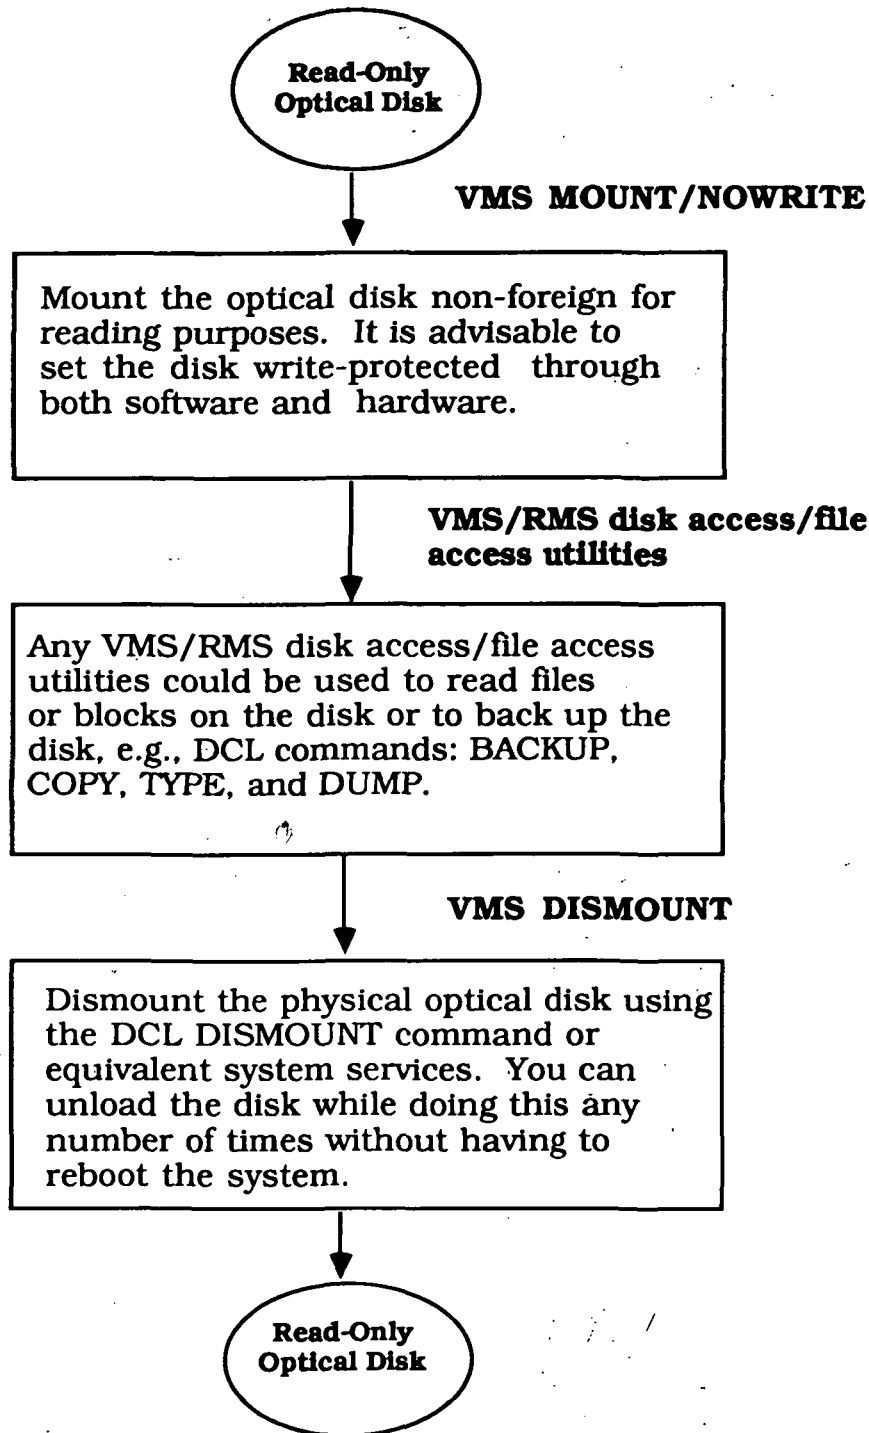


Figure 3. Actions taken during the read phase of the optical disk.

The QSDRIVER and SOAR software are no longer needed and should not be used. All operations should be performed on the physical optical disk.

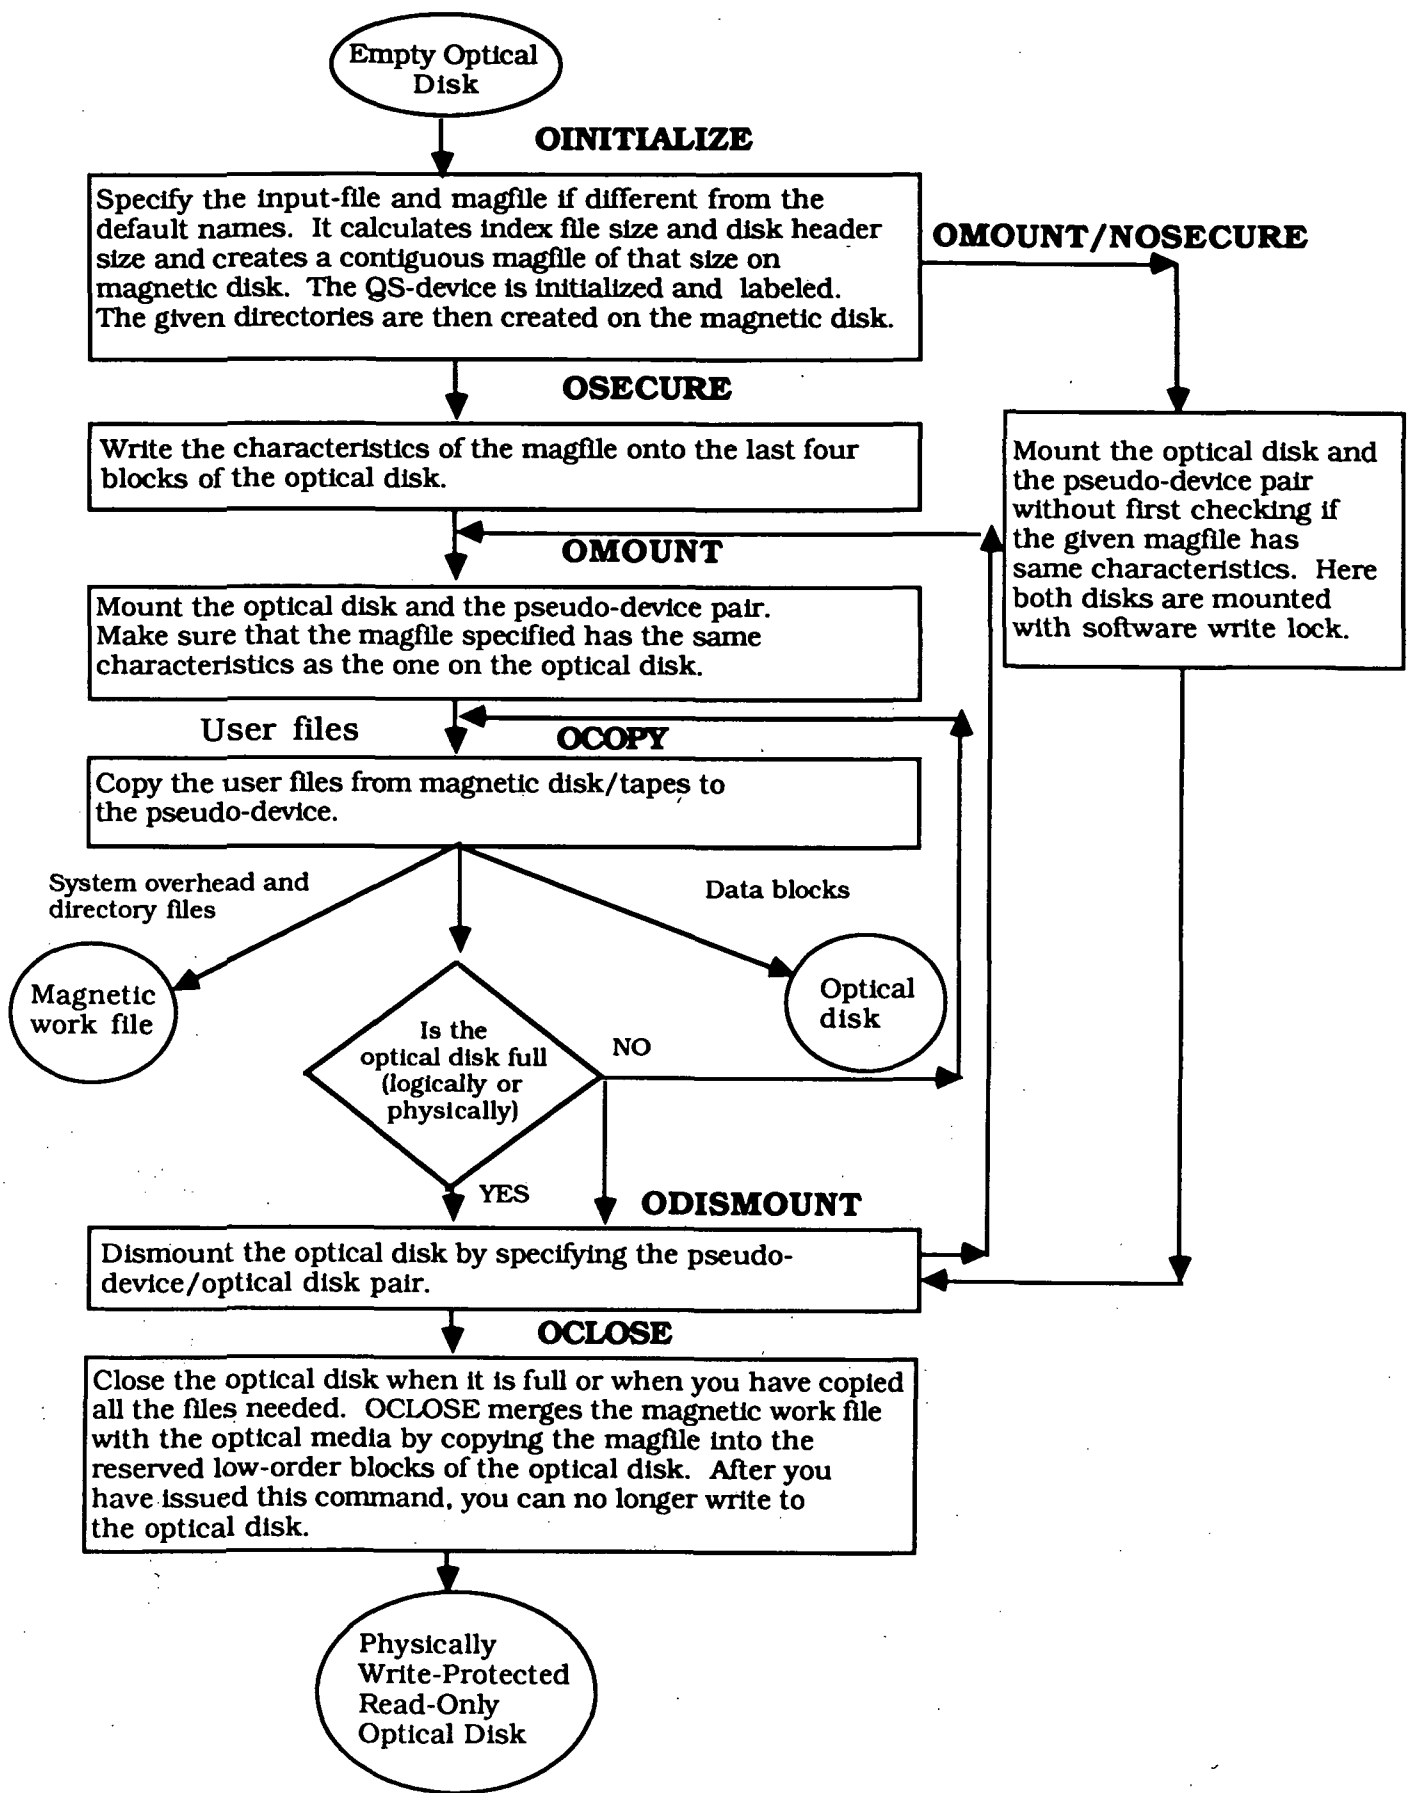


Figure 4. Actions taken during the write phase of the optical disk.

All SOAR operations are done on the QS-device. Only the OCOPY command should be used to write to the device.

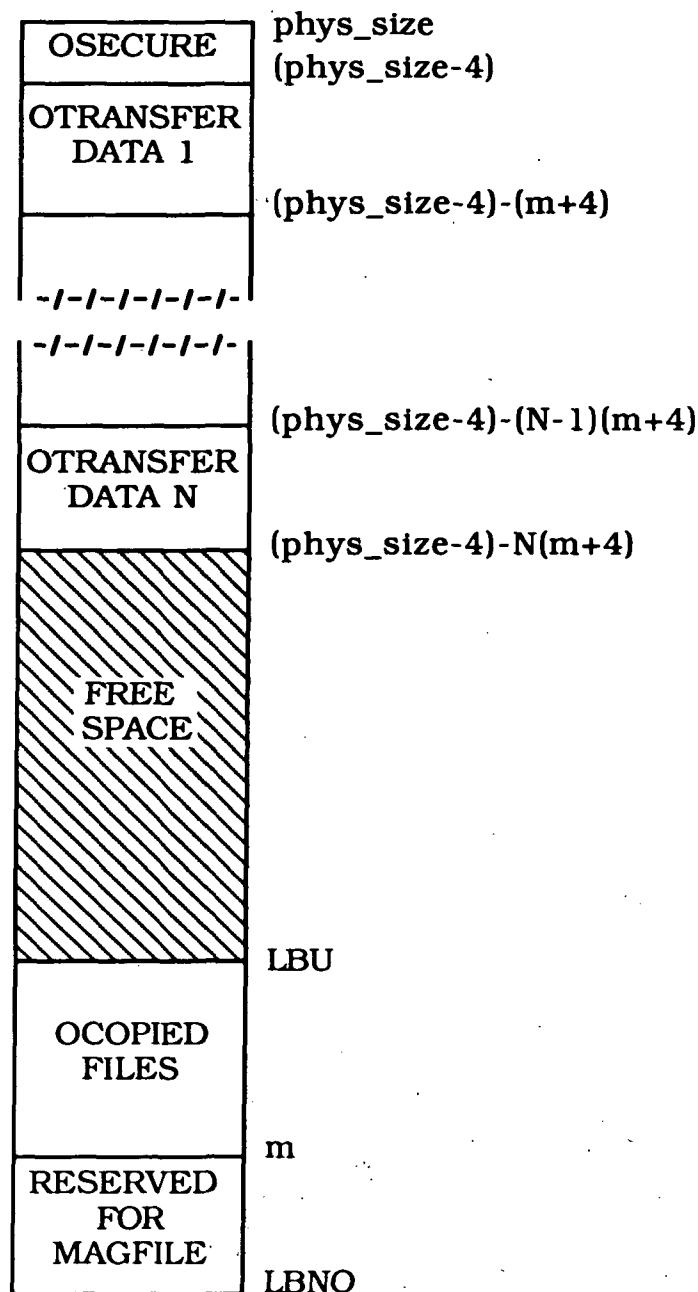


Figure 5. Optical disk configuration under SOAR.

LBNO = first block on the physical device

m = size of magfile in blocks

LBU = last block used by OCOPIED files

phys\_size = last physical block of device

phys\_size-4 = start of OSECURE data

(phys\_size-4)-(m+4) = start block of 1st OTRANSFER area\*

(phys\_size-4)-N(m+4) = start block of Nth OTRANSFER Area\*

\* Will not exist if OTRANSFER command is not issued. In this case the area will be free space available for OCOPY.

The most important phase of the optical disk life cycle is planning. Because the SOAR software writes files in VMS Files-11 format, it is necessary to keep the DEC Files-11 logical volume structure. Therefore, it is necessary to confine the system overhead files and the user directories to the low-order blocks of the SOAR logical disk. Once a file is written to a disk, any change to it (e.g., DCL commands DELETE, SET FILE, or SET PROTECTION) involves changing system overhead/directory information. Also, the DCL MOUNT and SET DEVICE commands write to the beginning blocks of the disk. To avoid corruption of the write-once optical disk and to be able to change file/disk characteristics during the write phase, these low-order blocks are kept in the magnetic file on a magnetic disk during the write phase. This magnetic file must exist (and hence its anticipated size must be known) before the OINITIALIZE module initializes the pseudo-device using the VMS INITIALIZE command. For this reason, the user must input the directory names and the number of files to be contained in each directory to the OINITIALIZE utility.

After initialization of the optical disk, the user must execute the OSECURE command, which writes magfile characteristics onto the last four blocks of the optical disk. This avoids corrupting the optical disk by the use of a wrong magfile by the user. After the disk is labeled by OSECURE, the OMOUNT and OCLOSE commands compare these characteristics against the characteristics of the specified magfile before mounting the device(s) successfully. Please note that after the execution of the OSECURE command, the OINITIALIZE command must not be executed during the life of the optical disk. However, before executing the OSECURE command, the user may mount the disk for read access to make sure that the directory tree structure is satisfactory by using the OMOUNT/ NOSECURE command. At this point, if the directory structure is not satisfactory, the OINITIALIZE command can be re-executed with the correct directory structure specifications. Once the OSECURE command has been executed, it is no longer possible to add directories. Files are copied to the disk using the OCOPY command and are copied from the disk using the standard VMS and RMS utilities (e.g., TYPE, PRINT, BACKUP, COPY, OPEN, and CLOSE).

**WARNING:** Any attempt to write to the disk using anything other than the OCOPY command could corrupt the files; the results are unpredictable.

Files may be virtually deleted from the optical disk using the DCL command SET FILE/REMOVE. This command deletes the directory entry for the file but leaves its file header in [0,0]INDEXF.SYS. Space on an optical disk can never be reused once it has been written to.

The final phase of the life cycle is entered when the user has decided that the disk is full. A single Files-11 formatted optical disk is produced by using the OCLOSE command. This command copies the magfile on the magnetic

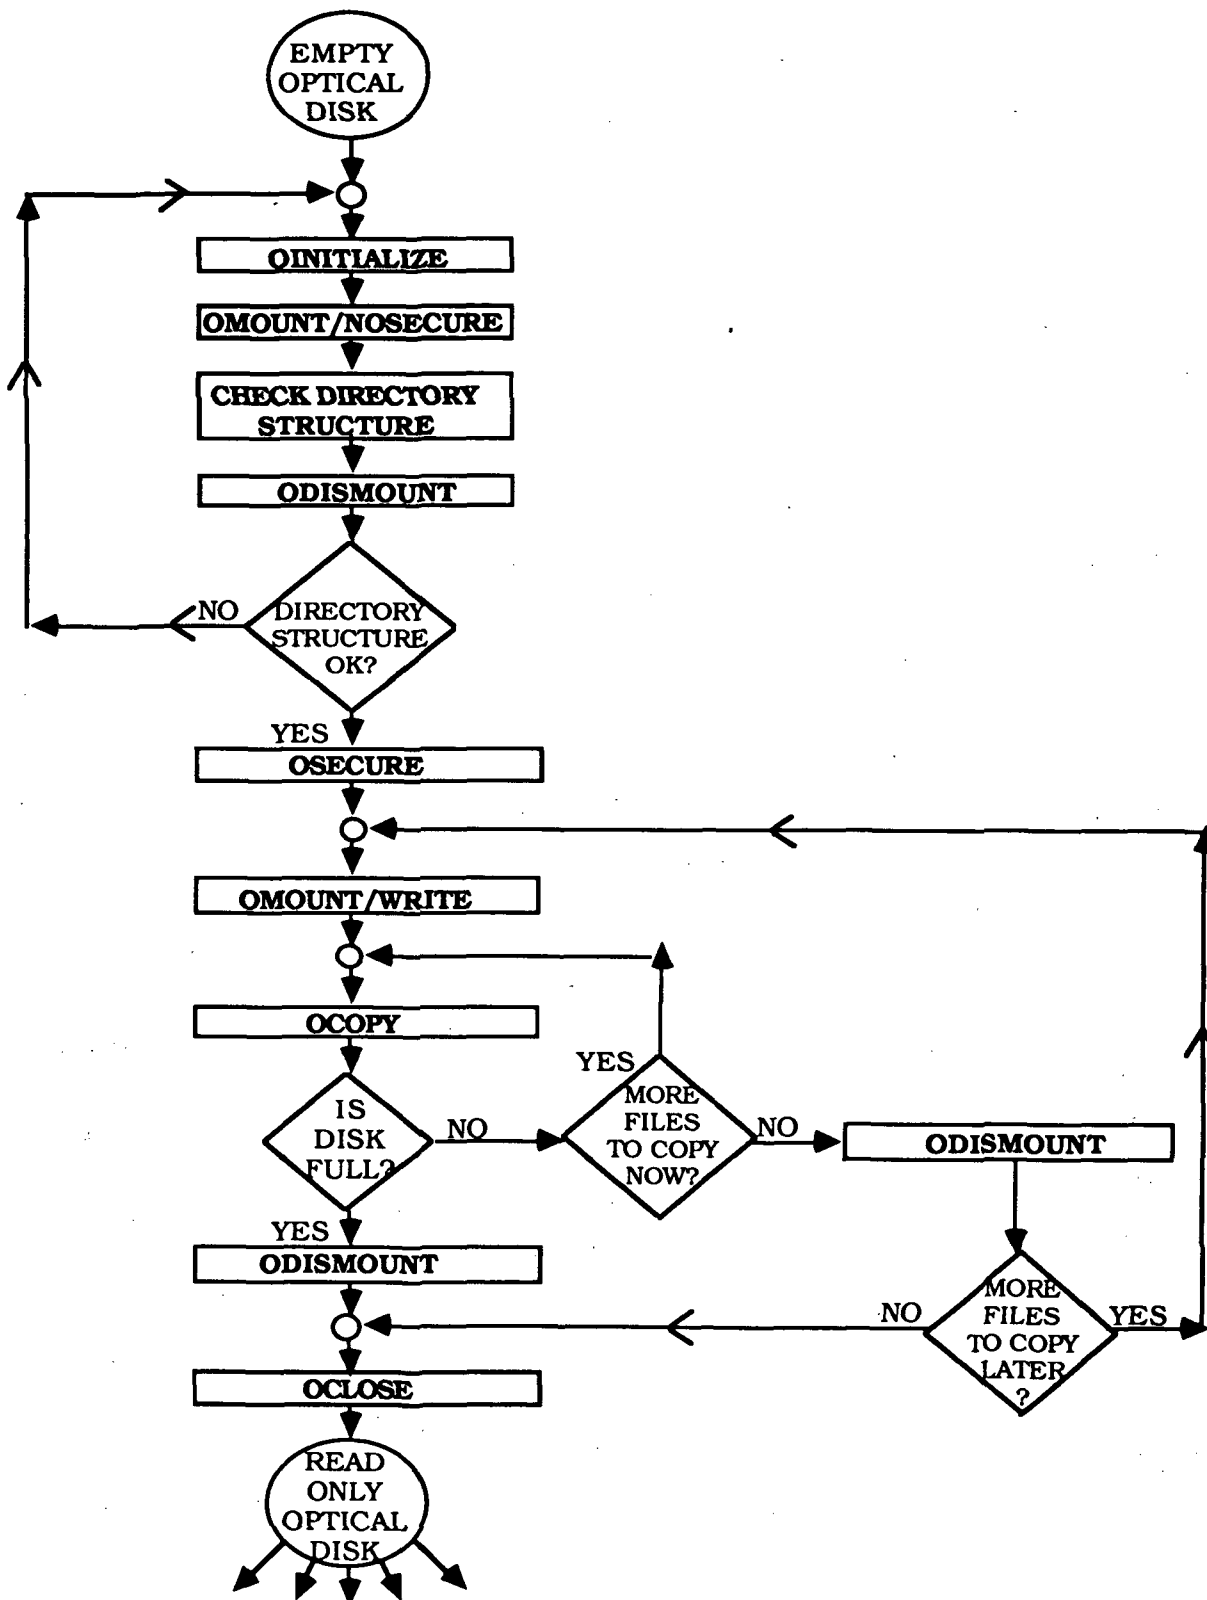


disk into the beginning blocks of the optical disk that were reserved during OINITIALIZE.

Once the OCLOSE operation has been completed, the optical disk should immediately be hardware write-locked. Mounting the closed optical disk via VMS "\$MOUNT" for write access (the default) will overwrite the BITMAP.SYS file and prevent subsequent disk mounts, rendering the disk useless. After the OCLOSE operation, the disk can be accessed using the standard VMS commands MOUNT/NOWRITE, COPY, BACKUP, etc. Neither the SOAR device driver nor the SOAR utility routines are used. In fact, the disk can be treated like any magnetic random read-only access disk. All the I/O methods to access files/records on VMS can be used for the files on this disk. No more write transfers can be done. Figure 6 shows the optical disk life cycle during read and write phases.

**Note:**

It is strongly recommended that the optical disk be both hardware and software write-locked at all times when data are not being written to the disk. As an extra precaution, files should be opened for read-only access.



VMS/RMS file access methods (see Figure 3)

Figure 6. Life cycle of the SOAR optical disk.



## 6.0

## The SOAR Command OCLOSE

OCLOSE copies the directory information from the magnetic disk work file (created by OINITIALIZE) into the reserved lower blocks of the optical disk. This operation is done only when all write transfers to the disk have been completed. **Immediately** following this command, physically write-protect the optical platter.

Privileges needed: LOG\_IO, VOLPRO.

### 6.1 Format

OCLOSE                pseudo\_device: label        optical\_device:

Qualifiers                                Default

/MAGFILE=<filename>                /MAGFILE=SOAR:label.\$SOAR\$

### 6.2 Parameters

pseudo\_device

The logical name that associates the QSDRIVER with the given optical disk drive (e.g., QSA0, QSA1,...).

label

Specifies the identification to be encoded on the optical disk. The label can consist of up to 12 alphanumeric characters.

optical\_device

Name of the optical disk drive.

### 6.3 Description

The OCLOSE utility is used after the user has decided that the optical disk is full or that it will not be written to any longer. OCLOSE merges the magnetic work file (created by OINITIALIZE) into the reserved lower blocks of the optical disk. Before doing this it mounts the optical device and then compares the specified magfile characteristics against the data written on the last four blocks of the optical disk by the OSECURE command. On failure of this command, all the devices mounted through this module are dismounted before the command aborts with an error message. Please note that, after execution of this command, the optical disk can be treated as a Files-11 formatted disk and can be accessed by using the standard VMS utilities (DCL, RMS, etc.). From then on, the user should not use SOAR to access the disk. Please refer to the section on Use of VAX VMS Commands for more details.

**WARNING:** This command must be used only once in the optical disk's life cycle (after the user has decided that the optical disk is full and it will not be written to any longer). After this command has been issued, the optical disk must **always** be mounted for **read-only** access. After the optical disk has been closed, the physical tab on the optical disk should be positioned to write-protect it before mounting. Failure to heed this warning may result in the corruption of the data on the platter.

#### **6.4 Qualifiers**

/MAGFILE=<filename>

The magnetic work file contains the system overhead files, user directory information, and header information of the files written to the pseudo-device. If the file name specification does not contain the disk and directory location, then the disk and directory pointed to by the SOAR logical name is used. If the directory location but not the device location is specified, then the current default disk is used.

#### **6.5 Examples**

\$ OCLOSE/MAGFILE=SOAR:OPTO.\$SOAR\$ QSA0: OPTO DUA0:

This example copies the magnetic work file (which contains the directory information) specified in the OINITIALIZE command onto the reserved lower blocks of the optical disk DUA0:.

## 7.0

## The SOAR Command **OCOPY**

OCOPY copies an input file from a magnetic disk or tape to an output file on the pseudo-device. The OCOPY command supports sequential, indexed, relative, and direct files.

**WARNING:** Do not use the DCL COPY command to copy the files because it may corrupt the optical disk; the results are unpredictable. Do not use DELETE!! If necessary, always use SET FILE/REMOVE.

Privileges needed: SYSPRV (needed to mark bad blocks).

### 7.1 Format

OCOPY	input_file	output_file
Qualifiers		Default
/CLUSTERSIZE=value		/NOCLUSTERSIZE
/DATACHECK		/DATACHECK
/ERROR		/NOERROR
/EXCLUDE=filename		/NOEXCLUDE
/LOG		/NOLOG
/VERIFY		/NOVERIFY

### 7.2 Parameters

#### input\_file

Specifies the location and name of the input file resident on the magnetic disk device.

#### output\_file

Specifies the location and name of the output file that will be written to the pseudo-device.

#### Note:

Output file version numbers are **always** maximized, even when a lower value is specified. Unless a higher version already exists or a different one is specified, OCOPY will preserve the version number. The input file creation date, owner UIC, and journal control flags are also preserved. The output file protection is always set for Read and Execute only, for all users (System, Owner, Group, and World). This is done to prevent accidental file deletion from the optical disk (see warning above). If you need to remove a file entry from a directory, be sure first to set the file protection for delete.

### 7.3 Description

The OCOPY command copies an input file from a magnetic disk or tape to the output file on the pseudo-device. As a result, the directory/file header information of the output file is written to the magnetic work file while the actual data are written to the optical disk. This operation does error checking while reading the input file. If an error is detected on Read, OCOPY will retry three times before aborting.

#### Note:

The output directory on the pseudo-device must be created during the initialization of the optical disk by using the OINITIALIZE command. This is done by including the directory name in the file specified by the /INPUT qualifier of the OINITIALIZE command.

### 7.4 Qualifiers

/CLUSTERSIZE=value

This qualifier is used to speed up the rate of the OCOPY transfers. Normally, OCOPY uses the cluster value established for the platter when it was OINITIALIZED. When OCOPY/CLUSTER is used, the file is extended by the value specified for each data transfer, thereby reducing the total number of I/O operations required. The cluster value must be an **even** integer between 2 and 124, inclusive. If the cluster value given here is less than the disk cluster value, then the disk cluster value is used. Allocation sizes of files OCOPIED using this qualifier are automatically adjusted to appropriately reflect the OINITIALIZE cluster value.

/DATACHECK

This qualifier specifies that a data check is to be performed after each write operation completes. Its use further ensures the integrity of the data being written; however, it will increase the amount of time it takes for OCOPY to complete. Data checks are done automatically when /VERIFY is selected. It is recommended that users **always** select data checks. The default is /DATACHECK.

/EXCLUDE=filename

This qualifier is used to exclude a file or group of files from the OCOPY operation. Wildcard file names are acceptable for this; however, only one file specification is permitted for exclusion.

## **/LOG**

This qualifier determines whether the file specification of each file processed is displayed at SYS\$OUTPUT during the operation. The default is /NOLOG. Note that the displayed information shows the actual number of blocks used by the output file, rather than the allocated number of blocks.

## **/VERIFY**

This qualifier specifies whether the software revectoring of the bad blocks should be enabled during the operation. The default is /NOVERIFY. The software revectoring requires that the hardware revectoring be enabled. The optical disk *should* be error free; that is, there should be no bad blocks on the disk. However, if bad blocks are detected and hardware revectoring of the bad blocks fails, the software attempts software revectoring, five times for each bad block encountered. After the fifth attempt, the software either aborts the OCOPY operation with an error message, or (if the /ERROR qualifier was given) continues with the next write operation.

### **Note:**

1. The privilege SYSPRV or volume ownership of the optical disk is required for the /VERIFY qualifier to detect bad blocks and revector them to a good spot on the disk.
2. Following a parity error, the OCOPY operation may not always be able to successfully software revector bad blocks. If the message

%OCOPY-F-VERIFY, verify operation does not work

is displayed, then retry the OCOPY operation.

## **/ERROR**

Normally, integrity of files copied is necessary for disk write operations. However, because of the write-once nature of the optical disk (once written, blocks cannot be overwritten), some users have shown an interest in being able to continue OCOPY even after bad blocks have been discovered in the output file. With this qualifier, OCOPY will continue after failure to revector bad blocks (hardware/software), to a maximum of 10 times.

## **7.5 Examples**

```
$ OCOPY/CLUSTER=50/EXCLUDE=*.COM *.C*   QSAO:[DATA]
```

In this example all files with a file type starting with the letter C, except for .COM files, are OCOPIED to QSAO:[DATA]. The files are OCOPIED 50 blocks at a time.



```
$ OCOPY/LOG  SYS$USER:[USER]IMAGE.DAT  QSA0:[DATA]
```

```
%OCOPY-S-COPIED, SYS$USER:[USER]IMAGE.DAT copied to  
QSA0:[DATA]IMAGE.DAT (4000 blocks)
```

The above example copies a 4000-block file from the SYS\$USER:[USER] directory to the directory [DATA] on the pseudo-device referenced by QSA0: (the data file is copied to the optical disk; the directory information about the file is stored on the magnetic work file).

```
$ OCOPY/VERIFY/ERROR SYS$USER:[USER]IMAGE.DAT  
QSA0:[DATA]*./LOG  
%OCOPY-W-DATACHECK, datacheck error occurred at VBN 1.  
%OCOPY-I-COPIED, IMAGE.DAT copied to QSA0:[DATA]IMAGE.DAT  
(4000 blocks)
```

In the above example, the file SYS\$USER:[USER]IMAGE.DAT was copied to QSA0:[DATA]IMAGE.DAT. However, this time a single bad cluster was found when trying to write the first virtual block of the file. Therefore, software revectoring of the bad blocks was enabled and the bad block was successfully revectoring within five tries. The copy was completed successfully. VBN refers to the Virtual Block Number on the optical disk.

```
$ OCOPY/VERIFY/ERROR SYS$USER:[USER]IMAGE.DAT QSA0:[DATA]*./LOG  
%OCOPY-W-DATACHECK, datacheck error occurred at VBN 497  
%OCOPY-I-VERIFY, software bad block revectoring unsuccessful 5 times.  
%OCOPY-W-ERROR, ignoring parity-datacheck error 1 times.  
%OCOPY-S-COPIED, SYS$USER:[USER]IMAGE.DAT copied to  
QSA0:[DATA]IMAGE.DAT (4000 blocks)
```

In the above example, the file SYS\$USER:[USER]IMAGE.DAT was copied to QSA0:[DATA]IMAGE.DAT. However, a single bad cluster was found during the operation and software revectoring of the bad blocks was tried. After five attempts, the bad block revectoring was unsuccessful; however, because the /ERROR qualifier was specified, the bad block was ignored and the OCOPY operation continued successfully.

**WARNING:** Do **not** use the VMS DELETE command to eliminate unwanted files from your optical platter. Doing so will free up blocks which have already been used. Use SET FILE/REMOVE instead.

## 8.0

## The SOAR Command **ODISMOUNT**

ODISMOUNT releases the given pseudo-device pair that was previously mounted using the OMOUNT utility.

Privileges needed: SYSNAM, only if the pseudo-device pair was mounted system-wide.

### 8.1 Format

ODISMOUNT    pseudo\_device:    optical\_device:

Qualifiers        Default

/CLUSTER	/NOCLUSTER
/NULL	/NONULL

### 8.2 Parameters

pseudo\_device

The logical name that associates the QSDRIVER with the given optical disk drive (e.g., QSA0, QSA1,...).

optical\_device

Specifies the name of the optical disk that needs to be dismounted and deallocated.

### 8.3 Description

ODISMOUNT dismounts and deallocates both the given optical disk and the pseudo-device that were mounted using the OMOUNT utility.

### 8.4 Qualifiers

/CLUSTER

This qualifier is used to dismount the volume clusterwide. After ODISMOUNT successfully ODISMOUNTS the volume on the local node, the volume is ODISMOUNTED on every other node in the existing VAXcluster.

#### Note:

The /CLUSTER qualifier may not operate on VMS versions earlier than 5.2.

## **/NULL**

This is used if the pseudo-device was OMOUNTED using the /NULL qualifier. In this case only the pseudo-device is dismounted. If /NULL is specified, then do not specify an optical device. Please see section 15.0, Use of the /NULL Qualifier, for more information.

### **8.5 Examples**

**\$ ODISMOUNT        QSAO:        DUAO:**

In this example the pseudo device QSAO:, and the associated optical device DUAO:, are both dismounted.

**\$ ODISMOUNT/NULL QSAO:**

In this example it is assumed that no optical device has been associated, so only the pseudo-device QSAO: is dismounted.

## 9.0

## The SOAR Command **OINITIALIZE**

**OINITIALIZE** calculates the total number of directory information and system overhead blocks (say *M*) needed to store the given number of directories/files on the optical disk and creates a magnetic work file of that size. Then it configures the pseudo-device, which consists of the magnetic work file and the optical disk (first *M* blocks reserved for directory information). Finally, it initializes the pseudo-device (using the VMS **INITIALIZE** utility) and creates directories, specified by the **/INPUT** file, on the pseudo-device.

**WARNING:** For certain controllers (such as the Emulex controller) it may be necessary to format the optical disk (a procedure that puts the Replacement Caching Table [RCT] into the beginning blocks of the disk) before you initialize it. Refer to the section on Command File for RCT Formatting, and to the appropriate sections of your controller manual.

Privileges needed: **LOG\_IO**, **VOLPRO**, **TMPMBX**.

### 9.1 Format

**OINITIALIZE**                      pseudo\_device: label      optical\_device

Qualifiers

Default

**/CLUSTER\_SIZE=nn**  
**/INPUT=<filename>**  
**/MAGFILE=<filename>**  
**/MAG\_SIZE**  
**/NULL[=<filename>]**

**/CLUSTER\_SIZE=8**  
**/INPUT=SOAR:label.INP**  
**/MAGFILE=SOAR:label.\$SOAR\$**  
**/NOMAG\_SIZE**  
**/NONULL**

### 9.2 Parameters

**pseudo\_device**

The logical name that associates the **QSDRIVER** with the given optical disk drive (e.g., **QSA0**, **QSA1**,...).

**label**

Specifies the identification to be encoded on the pseudo-device and on the optical disk; it is desirable to give each disk a unique label. VAX VMS Version 4.x limits the label length to a maximum of 12 alphanumeric characters.

**optical\_device**

Name of the optical disk drive. Identifies the physical optical disk to be initialized.

### 9.3 Description

The OINITIALIZE utility initializes the pseudo-device (QS-device) on which users can create their files (using the OCOPY command). It first checks the validity of values provided for all the parameters and qualifiers and gives default values when the user has not specified the values. Then it calculates the anticipated number of blocks on the disk to be taken up by system files and directory information (from the total number of directories/files given by the user) and creates the magfile of that size at the given location. Finally, after checking the physical device's characteristics, it configures the QSDRIVER to reflect the same device characteristics as the optical disk and to have correct pointers to the magfile, which is to be the first part of the pseudo-device. Then it initializes the pseudo-device, using the VMS INITIALIZE utility, and creates the specified directories on it.

#### Notes:

1. On failure of this command, all the devices mounted during its execution are dismounted before it aborts with an error message, and the optical disk is not written to.
2. The magnetic work file is of the utmost importance; great care should be taken to make sure that there is an adequate backup of the file at all times. The magnetic work file backup should be identified explicitly by name.

### 9.4 Qualifiers

/CLUSTER\_SIZE=nn

Every disk on VAX VMS has a characteristic parameter called CLUSTER\_SIZE, which dictates how many blocks on the disk should be allocated to a file when it is created. Allocated blocks of a file are the nearest multiple of the cluster size of the disk and are greater than or equal to the actual size of the file. If disk cluster size equals eight and file TTT.DAT contains 19 blocks of data, then DCL command \$ DIR/SIZE=ALL TTT.DAT displays the following information:

TTT.DAT	19/24
---------	-------

For the OCOPY operation, the data transfer rate is directly proportional to the chosen cluster size. However, if small files are being copied, a large cluster size wastes disk space. These two factors should be considered in choosing the cluster size of the optical disk.

**Note:**

The cluster size cannot be reset after a disk is OINITIALIZED. The cluster size specified will determine the cluster size for both the pseudo-device and the optical disk. Since the block size for an optical platter is 1024 bytes (vs. 512 for magnetic), the cluster\_size disk qualifier must be an **even** number between 2 and 124, inclusive. The default value for this qualifier is 8.

/INPUT=<filename>

This file describes what will be contained in the optical disk when all the files are written to it. In other words, it specifies the names of the directories and the number of files each directory or subdirectory contains. The format of each record in this file is

[directory\_name]/#\_files.

where #\_files is a positive integer specifying the number of files that will be included in that directory or subdirectory.

**Note:**

If a directory contains subdirectories, the record for the parent directory must precede the record for any of its subdirectories.

**Example:**

Suppose it is necessary to copy a directory tree [A...] where the directory A contains 20 files plus two subdirectories, B and C (each of which contains 30 files). Further, directory B contains a subdirectory D that contains 10 files. The input file (default name SOAR:SOAR\_DIR.INP) will look as follows:

```
[A]/20
[A.B]/30
[A.C]/30
[A.B.D]/10
```

Any parent directories must be created before the subdirectories.

It is highly recommended that once you compute the total number of directory entries to be used on your platter you increase this number by 50 percent. Whenever you remove a directory item or encounter an OCOPY error, a file header is used up in the index file. Increasing the number of directory entries will not significantly affect performance of your platter. In addition, it's a good idea to include at least a few spare directories.

**/MAGFILE=<filename>**

A magnetic work file with this name is created and initialized. Its size is determined by the entries in the input file. An exact amount of space is also reserved in the low-order blocks of the optical disk. This file contains the system overhead files and the user directories.

**/MAG\_SIZE**

This qualifier is used to print to the output device the amount of contiguous space the magfile will require. This is based on the cluster size specified and on the contents of the directory file specified with the /INPUT qualifier. After the magfile size is displayed, OINITIALIZE stops without actually creating or initializing a magfile. The use of this qualifier is strictly informational.

**/NULL[=filename]**

This qualifier is used to create a magfile when no physical device is readily available. The optional file name specifies a file containing a list of characteristics for the particular type of optical disk drive you will be using. If no file name is specified here, then default parameters for a 1 Gigabyte Optimem 1000 optical disk drive are used. If /NULL is used, then do not specify an optical device on the command line. Please see the section on Use of the /NULL Qualifier for more information.

## **9.5 Examples**

```
$ OINITIALIZE/INPUT=SOAR:OPTO.INP/MAG=SOAR:OPTO.$SOAR$-  
_ $ /NULL=DUAO.PAR QSAO: OPTO
```

This command creates a magfile called SOAR:OPTO.\$SOAR\$ and initializes the pseudo-device. Since the /NULL qualifier is used, no physical device is specified. The characteristics for the actual physical device are listed in DUAO.PAR.

```
$ OINITIALIZE/INPUT=SOAR:OPTO.INP QSAO: OPTO DUAO:
```

This creates a magnetic work file called SOAR:OPTO.\$SOAR\$ and initializes the pseudo-device. The qualifier /MAGFILE could have been used to give the magnetic work file a different name. The optical disk is labeled OPTO; it is advised that the optical disk and its plastic cartridge be physically marked with the same label. The cluster size of both the pseudo-device (QSAO:) and the optical disk (DUAO:) are set to 8.

```
$ OINITIALIZE/MAGFILE=HSC$DUA2:[SOAR]OPTO.INP-  
_$/CLUSTER_SIZE=124 QSAO: OPTO DUAO:
```

This creates a magnetic work file called DIRFILE.DES on the magnetic disk HSC\$DUA2 and initializes the pseudo-device. The file SOAR:OPTO.INP specifies how many directories will be copied to the optical disk, the directory names, the number of files each directory will contain, and the directory/subdirectory tree structure. The cluster size of both the pseudo-device (QSA0:) and the optical disk (DUA0:) is set to 124.

```
$ OINITIALIZE/CLUSTER_SIZE=32  QSA0:  OPTO  DUA0:
```

This creates a magnetic work file called SOAR:OPTO.\$SOAR\$ and initializes the pseudo-device. The information about names and the number of directories/subdirectories and files they contain should be in the file SOAR:OPTO.INP. The cluster size of both the pseudo-device (QSA0:) and the optical disk (DUA0:) is set to 32.

```
$  
OINITIALIZE/INPUT=SOAR:OPTO.INP/MAG=SOAR:OPTO.$SOAR$/CLUSTER=20-  
_$ /MAG_SIZE          QSA0:  OPTO  DUA0:
```

This example is intended to illustrate the use of the /MAG\_SIZE qualifier. When this command is issued, OINIT will print out a message stating the number of contiguous blocks required for the magfile SOAR:OPTO.\$SOAR\$. OINIT then terminates without creating the magfile. To actually create the magfile, this command should be issued again without the /MAG\_SIZE option.

Note:

In the above examples please notice the following:

/INPUT file name:	OPTO.INP
/MAGFILE name:	OPTO.\$SOAR\$
QS volume label:	OPTO

As a suggestion and as the examples illustrate, it is a good idea to keep the magfile name, the directory input file name, and the volume label the same. This way it is very easy to associate all of these items with one another. It is also good always to give the file type ".\$SOAR" to your magfiles and the file type ".INP" to the directory input files to see easily what files these are.





## 10.0

## The SOAR Command **OMOUNT**

First, the OMOUNT software validates the OMOUNT command parameters and qualifiers and provides default values wherever necessary. The software checks the file attributes of the magnetic work file, specified by the /MAGFILE qualifier, against the information written to the last four blocks of the optical disk by the OSECURE command, mounts the optical disk (/FOREIGN), and mounts the pseudo-device. Finally, the pseudo-device is configured to reflect the characteristics of the physical device and to point to the magnetic work file location.

Privileges needed: LOG\_IO and SYSNAM for OMOUNT/SYSTEM.

### 10.1 Format

OMOUNT pseudo\_device: label optical\_device:

Qualifiers	Default
/CLUSTER	/NOCLUSTER
/MAGFILE=<filename>	/MAGFILE=SOAR:label.\$SOAR\$
/NOSECURE	/SECURE
/NULL[=<filename>]	/NONULL
/OTRANSFER	/NOTRANSFER
/SYSTEM	/PROCESS
/WRITE	/NOWRITE

### 10.2 Parameters

pseudo\_device

Specifies the name of the pseudo-device and the optical disk (the one that has the given label) that it should mount.

label

Specifies the identification to be encoded on the optical disk. The label can consist of up to 12 alphanumeric characters.

optical\_device

Name of the optical disk drive.

### 10.3 Description

OMOUNT first matches the validity of values given for all the parameters and qualifiers and provides default values wherever necessary. Then it matches the file attributes of the magfile specified by the /MAGFILE qualifier against the information on the last four blocks of the optical disk. If the attributes

do not match, the command aborts with an appropriate error message. Otherwise, it mounts the optical disk (/FOREIGN) and the pseudo-device with the given label and configures the pseudo-device to reflect the same characteristics as the physical device, including pointers to the magfile location. If OMOUNT fails at any point during the operation, all the devices mounted by the module are dismounted before it aborts with an error message.

**Note:**

In spite of the OSECURE command, the magfile can still be relocated from one disk to another during the write phase of the optical disk by using the DCL commands COPY and BACKUP. However, no other commands must be used for this purpose.

## **10.4 Qualifiers**

### **/CLUSTER**

This qualifier is used to mount the volume clusterwide. After OMOUNT successfully mounts the volume on the local node, the volume is mounted on every other node of the existing VAXcluster. After using this qualifier, use ODISMOUNT/CLUSTER to dismount the volume.

**Note:**

The /CLUSTER qualifier may not operate properly on VMS versions earlier than 5.2.

### **/MAGFILE=<filename>**

A magnetic work file with this name is created and initialized by the OINITIALIZE command. When the command OCOPY is given, the system overhead information and the user directory information about the file are written to this file. The OCLOSE command then copies this file to the optical disk.

**Note:**

It is essential to give the correct magnetic work file name (i.e., the same file that was created during OINITIALIZE).

### **/NOSECURE**

With this qualifier, OMOUNT does not match the file attributes of the magfile specified by the /MAGFILE qualifier against the information on the last four blocks of the optical disk. The user may wish to use this qualifier after the OINITIALIZE command, but before the OSECURE command, to

verify that the directory tree created during OINITIALIZE is satisfactory. If not, OINITIALIZE can be re-executed to correct the problem.

**Note:**

With this qualifier, both devices (physical and pseudo) will be mounted software write-locked (the /WRITE qualifier will be ignored). This safety precaution prevents the user from executing the OMOUNT/NOSECURE command, overriding the magfile validation, and inadvertently writing to the disk when an incorrect magnetic work file has been specified.

**/NULL[=<filename>]**

This qualifier is used to OMOUNT the pseudo-device without associating it with any physical device. The optional file name specifies a file containing a list of parameters for the particular type of optical disk drive you normally will be using. If no file name is specified here, then default parameters for a 1 Gigabyte Optimum 1000 optical disk drive are used. If /NULL is used, then do not specify an optical device on the command line. The pseudo-device will be allocated and mounted read only. The qualifier is used to verify the directory structure of a magfile created using OINIT/NULL. It is also used for checking the directory of an older, previously written to magfile without having to actually load an optical platter.

ODISMOUNT/NULL should be specified afterwards. Since there is no optical device associated with the pseudo-device, the only valid disk function when /NULL is used is DIR. Please see the section on Use of the /NULL Qualifier for more information.

**/OTRANSFER**

With this qualifier, OMOUNT will mount the device using the magfile that was previously OTRANSFERRED to the end of the optical platter. The device is mounted software write-locked (the /WRITE qualifier is ignored). If no magfile is found, OMOUNT exits with an error message.

**/SYSTEM**

This qualifier allows the optical disk to be mounted for system-wide access.

**Note:**

The privilege SYSPRV or OPER is needed for this qualifier.

**/WRITE**

This qualifier allows the user to write to the mounted device. The user can then copy the files to the optical disk using the OCOPY utility.

**Note:**

After a user issues the OCLOSE command and signifies that the optical disk is full, the device should not be mounted with this qualifier; this procedure will corrupt the optical disk. Also, after mounting the device with the qualifier, an attempt to write to the disk with any command other than OCOPY could corrupt the file, or even the whole optical disk. It could also cause the computer to crash.

**10.5 Examples**

**\$ OMOUNT/NULL QSA0: OPTO**

In this example, QSA0: is OMOUNTED using SOAR:OPTO.\$SOAR\$ as the magfile. The device is allocated and mounted read-only. Notice that no optical device is specified here.

**\$ OMOUNT/SYSTEM/WRITE QSA0: OPTO DUA0:  
%OMOUNT-I-MOUNTED, OPTO mounted on QSA0:**

The optical disk labeled OPTO is mounted on pseudo-device QSA0: using the magnetic work file SOAR:OPTO.\$SOAR\$, after matching the file attributes of this file against information written in the last four blocks of the optical disk (through the OSECURE command). The device is mounted for system-wide access, and files may be copied to it by using the OCOPY utility.

**\$ OMOUNT/NOSECURE QSA0: OPTO  
DUA0:/MAGFILE=SOAR:OPTO.\$SOAR\$  
%OMOUNT-I-MOUNTED, OPTO mounted on QSA0:**

The optical disk labeled OPTO is mounted on pseudo-device QSA0: using the magnetic work file SOAR:OPTO.\$SOAR\$. Here the file attributes of this file are not checked against the information in the last four blocks of the optical disk. The device is mounted for process-wide access only and will not be available to other users. In addition, it is mounted for read access only. An attempt to write to it will generate an error message.

## 11.0

## The SOAR Command OSECURE

The OSECURE command writes information about the given magfile onto the last four blocks of the optical disk; therefore, the last four blocks of the optical disk cannot be written to using the OCOPY operation. After the execution of this command, every OMOUNT operation will match this information on the optical disk with the characteristics of the given magfile.

### Note:

This command may be executed only once during the life of the optical disk. After execution of this command, the OSECURE command should not be attempted again.

Privileges needed: LOG\_IO, VOLPRO.

### 11.1 Format

OSECURE            pseudo\_device:    label        physical\_device:

Qualifiers

Default

/MAGFILE=<filename>

/MAGFILE=SOAR:label.\$SOAR\$

### 11.2 Parameters

pseudo\_device

The logical name that associates the QSDRIVER with the given optical disk drive (e.g., QSA0, QSA1, ...). This device is mounted by the OSECURE operation.

label

Specifies the identification to be encoded on the optical disk. The label can consist of up to 12 alphanumeric characters.

physical\_device

The physical device (optical disk) onto which the magfile attribute information will be written. This device is mounted /FOREIGN during the operation and must be hardware write-enabled.

magfile

Specifies the location of the magfile that contains the directory information about the given physical device. Here you must give the specification of the same magfile that was created during the OINITIALIZE operation.

File attributes of this file are written into the last four blocks of the given optical disk during this operation. The attributes written are as follows:

creation_date	file_name and type
file_size (actual)	file_organization
file_attributes	record_format
record_attribute	

### 11.3 Description

The OSECURE command writes information about the given magfile onto the last four blocks of the optical disk and creates the file [000000]OSECURE.SYS on the optical disk. After execution of this command, the OMOUNT and OCLOSE operations will compare this information with the characteristics of the specified magfile to ensure that the correct magfile is used with the optical disk.

### 11.4 Qualifiers

/MAGFILE=<filename>

This qualifier specifies the name of the magnetic work file that will be permanently associated with this optical disk. If this qualifier is not specified, then the default file name SOAR:label.\$SOAR\$ will be used.

### 11.5 Example

\$ OSECURE/MAGFILE=SOAR:OPTO.\$SOAR QSA0: OPTO DUA0:

In this example, QSA0: is mounted for write. It is then searched for the file QSA0:[000000]OSECURE.SYS and, if located, OSECURE aborts. If the file is not found, it is created and the attributes of the file SOAR:OPTO.\$SOAR\$ are written to the last four blocks of DUA0:. QSA0: is then dismounted.

## 12.0

## The SOAR Command **OSTATUS**

The **OSTATUS** command mounts and reads the specified optical platter and reports the status of that platter. It will inform the user if the platter is closed and, if not, it will identify the name, date, and size of the associated magfile and will return the dates, if any, of the latest **OTRANSFER** operations. It then dismounts the platter.

Privileges needed: None.

### 12.1 Format

**OSTATUS**                      **Physical\_device**

**Qualifiers**                      **Default**

**/MAG**                              **/NOMAG**

### 12.2 Parameters

**Physical\_device**

The physical device that contains the platter to be checked.

### 12.3 Description

The **OSTATUS** command first mounts the device **/FOREIGN/NOWRITE**. It next checks if there are any data present in the reserved magfile area at the beginning of the platter. If so, it prints a message stating that the platter is closed and then dismounts and exits. Otherwise, it next reads the **OSECURE** data stored in the last four blocks and prints out the associated magfile name, creation date, and size in blocks. Then it checks for the presence of any **OTRANSFER** data and, if found, prints the date of the latest **OTRANSFER** to the platter. It next prints the date, if any, of the **OTRANSFER** from the platter and then dismounts the platter and exits.

#### **Note:**

The **OTRANSFER** dates inform the user whether any **OTRANSFERS** have occurred and, if so, when. If there is no **OTRANSFER** from-date, then the **OTRANSFER** data *may* be consistent with the disk structure. They may not be consistent, however, if more files were **OCOPIED** to the disk after the **OTRANSFER** operation. If there is an **OTRANSFER** from-date, then it is less likely that the **OTRANSFER** data on disk are an accurate representation of the disk structure, because someone has taken the time to restore the magfile back onto magdisk with the probable intention of using it. For this reason, secure knowledge of the magfile and **OTRANSFER** use history must be known by the user. Also, note that the disk can be **OMOUNTED** for read-only using the directory information contained in the **OTRANSFER** area. It



is not necessary to OTRANSFER back from optical platter to magfile in order to mount the disk. Please refer to section 14.0, The SOAR Command OTRANSFER, and section 11.0, The SOAR Command OMOUNT, for more information.

#### **12.4 Qualifiers**

This qualifier is used during magnetic emulation.

#### **12.5 Example**

OSTATUS

DUA1:

In this example, the platter in drive DUA1: is mounted /FOREIGN/NOWRITE and is tested for being closed. If not closed, the name of the associated magfile and any OTRANSFER dates are printed. The device is then dismounted.

OSTATUS/MAG

DUA1:

This example is just like the preceding one except DUA1: is a magnetic disk.

#### **Note:**

When OTRANSFER is used under MAG emulation, the OTRANSFER-FROM date is not written. Therefore, when using OSTATUS/MAG the OTRANSFER-FROM field output is unpredictable.

## 13.0

## The SOAR Command **OTRANSFER**

The OTRANSFER command copies the magfile from a magnetic disk onto the optical platter or from the optical platter back onto a magnetic disk, depending upon the order of the command arguments. This command allows the user to easily transport a platter with its directory information file without having to close the platter. Each time the directory information is OTRANSFERRED to the platter, blocks of the platter that would have been available for data will be used to hold the directory information. If OTRANSFERRING to the optical platter, it must be OMOUNTED/WRITE. If OTRANSFERRING from the optical platter, then it must not be mounted.

Privileges needed: LOG\_IO, VOLPRO.

### 13.1 Format

OTRANSFER            input            output

### 13.2 Parameters

input

Specifies the name of the magfile or physical\_device.

output

Specifies the name of the magfile or pseudo\_device.

### 13.3 Description

Before any data are written anywhere, the attributes of the specified magfile are checked against the OSECURE data stored at the end of the disk. If they match, the command processing continues; otherwise, an error message is printed and the command aborts. If the magfile is specified as input and the pseudo-device is specified as output, the OTRANSFER command writes all the data from the magfile to the end of the available data area of the optical platter. They are written immediately before the OSECURE data and any previously OTRANSFERRED magfiles. If the physical-device is specified as input and the magfile as output, then the optical platter is searched for the last OTRANSFERRED magfile and, if present, the magfile is created and the magfile data are transferred into it. The magfile is created with the same file attributes that are stored in the OSECURE data area. If no OTRANSFER data are found, then an error message is generated and the command aborts.

#### Note:

OTRANSFER may be executed as many times as the user wishes, until the disk is filled. Large magfiles will use a lot of space at the end of the disk; therefore, the OTRANSFER command should be used sparingly. The

intended use of OTRANSFER is to facilitate transportation of optical platters between SOAR users. Although OTRANSFER can be used to back up the magfile onto optical disk, to conserve optical disk space periodic backups of the magfile onto magnetic disk or tape are recommended instead. In addition, it is possible to OMOUNT the drive (read-only) using the data contained in the OTRANSFER area. Please see section on The SOAR Command OSTATUS and the section on The SOAR Command OMOUNT for more information.

Also, when OTRANSFERRING the magfile to optical disk, the file [000000]OTRANSFER.SYS is created and extended. Its size is given by the size of the magfile plus four additional blocks for the OTRANSFER dates.

### **13.4 Qualifiers**

None.

### **13.5 Examples**

```
$ OTRANSFER SOAR:OPTO.$SOAR$ QSA0:
```

In this example, the attributes of the magfile SOAR:OPTO.\$SOAR\$ are checked against the OSECURE data and, if they match, the magfile is transferred to the end of the free data area on the platter in QSA0:.

```
$ OTRANSFER DUB4: SOAR:OPTO.$SOAR$
```

The magfile SOAR:OPTO.\$SOAR\$ is created from the OTRANSFER data previously stored at the end of the the optical platter in DUB4:. If the magfile name does not match the one found in the OSECURE data area, the new magfile is not created.

### **Note:**

When OTRANSFER is used under MAG emulation, the OTRANSFER-FROM date is not written. Therefore, when using OSTATUS/MAG the OTRANSFER-FROM field output is unpredictable.

## 14.0

## Use of the /NULL Qualifier

There are times when the physical device is unavailable or inconvenient to use. Since certain operations do not do any reading or writing to the optical platter, the physical device is not really necessary; in these situations a new qualifier, /NULL, can be used. It is available for use with the following commands: OINITIALIZE, OMOUNT, and ODISMOUNT.

The /NULL qualifier should be used when you wish to OINITIALIZE a platter when an optical drive is unavailable, such as in certain jukebox environments. OMOUNT/NULL can then be used to verify the correctness of the magfile directory structure in much the same way that OMOUNT/NOSECURE is used. ODISMOUNT/NULL is used afterwards.

When OINITIALIZE/NULL or OMOUNT/NULL is selected, default values are used for the physical characteristics of the disk drive. These default values are as follows:

physical size	=	1998948
number of sectors	=	50
number of tracks	=	1
number of cylinders	=	39979

These values are standard for 1 GB Optimem drives (including Series-M) and platters. If these values are not acceptable for your application, then the /NULL qualifier should be specified with an input file name like this:

OINITIALIZE/NULL=PARAMS.LIST

The file PARAMS.LIST contains values for the items listed above for your particular application. Each item in this file should be entered on a separate line and only the value should appear. For instance, PARAMS.LIST might contain the following items:

```
138672
18
8
1020
```

where 138672 is the physical size in blocks, 18 is the number of sectors, 8 is the number of tracks, and 1020 is the number of cylinders. If these values are not known, then the supplied command procedure CREATE\_NULL\_LIST.COM can determine them for you and generate the necessary null parameters file. The physical device must be available in order for this to work. Once the NULL parameters file is created, it can be reused for subsequent OINITIALIZATIONS and OMOUNTS. To run the command file you must specify the physical device name and an optional output filename. The device characteristics will be written to the output file or, if no file is specified, to the file <devname>\_NULL.LIST. The file

<devname>\_NULL.LIST must then be used with the /NULL qualifier (/NULL=<devname>\_NULL.LIST) when you OINITIALIZE the magfile. Examples of the use of the command follow.

Example 1: \$ @CREATE\_NULL\_LIST DUC2:

This writes the necessary values to DUC2\_NULL.LIST.

Example 2: \$ @CREATE\_NULL\_LIST DUC2: DUC2.PARAMS

This writes the necessary values to DUC2.PARAMS.

**Note:**

The parameter list file can be used with any disk drive of the same type as the one specified when the file was created; i.e., DUC2\_NULL\_LIST can be used when OINITIALIZING or OMOUNTING any other DUC disk, as long as the drives are the same type.

Once the magfile has been created via OINIT or OINIT/NULL, it can be verified without using the optical device by issuing the OMOUNT/NULL [=<filename>] command. This command can be issued at any time during the life cycle of the magfile to view the SOAR disk directory. OMOUNT/NULL mounts the pseudo-device as allocated and read-only. ODISMOUNT/NULL is used to dismount the pseudo-device when OMOUNT/NULL is used. In all situations, OINIT/NULL, OMOUNT/NULL, and ODISMOUNT/NULL, no physical device should be specified; otherwise, an error message will be displayed.

## 15.0

## Use of VAX VMS Commands

Because the SOAR software writes data to the optical disks using standard DEC Files-11 format, a user can access (read- and execute-only) these data by using any of the VMS file-access methods (i.e., DCL commands or RMS utilities) during the life of the optical disk.

The following DCL commands can be used to read the files (written by using OCOPY) from an optical disk.

- To copy a file from the optical disk to magnetic tape or disk, or to another optical disk:

```
$ BACKUP QSA0:[DIRA]*.DAT MFA0:SAVE.BCK/SAVE_SET
$ BACKUP/LOG DUB4:[DIRA]*.DAT SYS$USER1:[USER]
$ COPY QSA0:[DIRA]*.DAT DUB4:[DIRA]
```

- To read, print, or compare files:

```
$ TYPE QSA0:[DIRA]README.DAT
$ PRINT QSA0:[DIRA]README.DAT
$ DIF QSA0:[DIRA]README.DAT DISK1:[USER]README.DAT
```

- To verify the integrity of the file:

```
$ DUMP/RECORDS QSA0:[DIRA]README.DAT
$ ANALYZE/RMS QSA0:[DIRA]README.DAT
$ DIRECTORY/FULL QSA0:[DIRA]README.DAT
```

There are also some DCL commands that can be used only during the write phase of the optical disk. These commands must not be used once the disk has been closed using the OCLOSE command. These commands change the characteristics of the files/QS-device.

- To change the characteristics of the file or to (virtually) delete a file:

```
$ RENAME QSA0:[DIRA]README.DAT QSA0:[DIRA]READ_ME.DAT
$ SET FILE/REMOVE QSA0:[DIRA]README.DAT (to delete the file)
$ SET FILE QSA0:[DIRA]README.DAT/PROTECTION=(WORLD:RE)
```

or

```
$ SET PROTECTION QSA0:[DIRA]README.DAT/-
_$ PROTECTION=(WORLD:RE)
$ SET FILE/OWNER=[73,73] QSA0:[DIRA]README.DAT
```

- To change the characteristics of the QS-device:

```
$ SET VOLUME/OWNER=[73,73] QSA0:
$ SET VOLUME/LABEL="YOURDISK1" QSA0:
```

Apart from DCL commands, file level and record I/O may also be executed by using standard RMS procedures or high-level language I/O procedures. It is strongly recommended that the optical disk be hardware and software write-locked at all times when data are not being written to the disk. As an extra precaution, files should be opened for read-only access.

There are some VMS functions you must **never** perform on your SOAR platter:

- Never issue the standard VMS MOUNT command unless you are **absolutely certain** that the OCLOSED platter you wish to mount is hardware write-protected. Otherwise, you **will** corrupt the optical platter.
- Never use the DELETE command to eliminate files from your SOAR directory. Always use SET FILE/REMOVE. Otherwise, you **will** corrupt blocks and possibly good files on your platter.

## 16.0

## Magnetic Emulation Mode

As with any other software, there may be initial difficulty in using the SOAR software on a given system. The main cause of this difficulty is usually a natural reluctance to use the software on the write-once optical disks, which cannot be erased once written to.

This problem can be avoided if you have access to a magnetic disk that can be initialized (i.e., one that you can erase). The disk can be used to emulate the optical disk environment using SOAR software. Once you are comfortable with the software, try using it on the optical disk. In this way, you can practice without worrying about making mistakes on a write-once medium. This will also help you later in distinguishing between hardware and software problems.

### How to Use SOAR Software in the Magnetic Emulation Mode

Read this user's guide thoroughly before using the software. In all the SOAR commands explained earlier in the document, wherever "optical\_device:" is asked for, it should be replaced with the name of the magnetic disk that will be used to emulate an optical disk. Please note that you should not choose the location of the magnetic file (or the location of any file needed by the software) on this disk.

For example, if

magnetic emulation disk is:	HSC\$DUA2:
pseudo_device is:	QSA0:
label is:	TEST
magfile is located at:	HSC\$DUA3:[SOAR]TEST.\$SOAR\$
input_file is located at:	HSC\$DUA3:[SOAR]TEST.INP

the SOAR commands would look like the following:

```
$ OINITIALIZE/INPUT=HSC$DUA3:[SOAR]TEST.INP/-  
_$ MAGFILE=HSC$DUA3:[SOAR]TEST.$SOAR$ QSA0: TEST HSC$DUA2:  
  
$ OSECURE HSC$DUA2: HSC$DUA3:[SOAR]TEST.$SOAR$  
  
$ OMOUNT/WRITE QSA0: TEST HSC$DUA2:/-  
_$ MAGFILE=HSC$DUA3:[SOAR]TEST.$SOAR$  
  
$ OCOPY/MAG input_file QSA0:[.....  
  
$ ODISMOUNT QSA0: HSC$DUA2:  
  
$ OCLOSE/MAGFILE=HSC$DUA3:[SOAR]TEST.$SOAR$ QSA0: TEST-  
_$HSC$DUA2:
```



**\$ OSTATUS DUA2:**

**\$ OTRANSFER HSC\$DUA3:[SOAR] TEST.\$SOAR\$ QSA0:**

## 17.0

## Command File for RCT Formatting

A command file is available that simplifies the otherwise cumbersome task of RCT initialization for users who have the prerequisites detailed below. The command file is not fully automatic, but with user-friendly prompts it leads the user through the whole process.

The prerequisites for using the procedure are the following:

Emulex UC14 or UC04 with firmware version E or below  
IBM PC  
PROCOMM program

The PC must be connected to the operator console port on the VAX. The PROCOMM is used for terminal emulation (VT-100) and for the execution of the command file.

The command file is executed interactively because the results of the EXAMINE commands must be verified visually by the operator before the next step is executed.

For RCT initialization, the following procedure must be followed by a privileged operator:

1. Shut down the system.

```
$ SET DEF SYS$SYSTEM  
$ @SHUTDOWN
```

When the shutdown is finished, hit the HALT button. The prompt >>> should appear on the screen.

2. Insert the optical platter to be initialized into the drive. The platter must not be write-protected.
3. Hit the START button on the drive and wait until the flashing green light on the drive is lit continuously.
4. Start the RCT.CMD command procedure.
  - a. Hit ALT/F5 on the keyboard.
  - b. Enter the command file name RCT.CMD or move the cursor to RCT.CMD on the menu.
  - c. Follow the prompts of the command procedure.
  - d. The last message of the procedure informs the operator about the success or failure of the initialization.
5. Remove the disk from the drive.

6. If more disks have to be initialized, repeat steps 2 through 5.
7. Reboot the VAX system.

## 18.0

## Software Installation

The software is distributed in a BACKUP save set on 1600-bpi magnetic tape. Note that only the object code is available for distribution. The procedure for installing the SOAR software from tape is as follows.

First, log onto the **SYSTEM** account and then mount the distribution tape foreign.

```
$ MOUNT/FOREIGN MFA0:
```

Next, copy the files from the save set on tape to the subdirectory [SYSMGR.SOAR].

```
$ BACKUP/VERIFY MFA0:SOAR.BCK  
SYS$SYSROOT:[SYSMGR.SOAR]*.*
```

The driver and utility routines must be linked on your system by using the **SOAR\_BUILD.COM** command procedure.

```
$ SET DEFAULT SYS$SYSROOT:[SYSMGR.SOAR]  
$ @SOAR_BUILD
```

At this point, a warning may appear:

```
LINK-W-USRTFR, image  
SYS$SYSROOT:[SYSMGR.SOAR]QSDRIVER.EXE has no user transfer  
address.
```

Please ignore this message.

To use the software, add the following lines to the system startup command file.

```
$ @SYS$SYSROOT:[SYSMGR.SOAR]SOAR_STARTUP
```

This command procedure presumes that the SOAR software is in **SYS\$SYSROOT:[SYSMGR.SOAR]**. If you have kept the software in another location, you will have to change this procedure accordingly. Also, once the optical disk is initialized as described in this user's guide, add the **OMOUNT** command. This will ensure that the optical disk is mounted properly.

There is also a command file, **LLOGIN.COM**, that will give ideas for creating symbols for the SOAR commands. To reduce typing effort, it can be used after making changes appropriate for the specific system. It will also ensure that the correct qualifier values are given. You may need to have one **LLOGIN.COM** for each pseudo-device on your system.



The following is a detailed explanation of the different jumper/cable settings involved in the Emulex/Optimem installations. Figures 7 and 8 show the daisy-chain layout for different Emulex configurations. Note that the dip switches shown in the figures are for a UC14 board installed on a Unibus at the CSR address 772150 (factory default value). The similar settings for UC04 installed on a Qbus at the same CSR address are shown at the bottom of the figures. Refer to the Emulex manual if you are using a different CSR address. Apart from the Emulex board setting, the Optimem connections should be the same for both the UC04 and the UC14.

There are four possible options with which to configure Optimem drives with an Emulex host adapter: 40, 41, 42, and 43. The first two options are available if you are planning to connect only the Optimem drives to UC04/UC14. However, if you are connecting a combination of optical and magnetic drives, you must choose options 42 or 43.

You can daisy-chain your Optimem drives in different ways to have a combination of master/slave drives. All the drives that are connected through the SCSI port are considered master drives, while the drives connected through the ODI port are considered slave drives. The slave drives share the controller and EDAC of their master drive; therefore, these boards *must* be removed from the slave drives.

SCSI to SCSI or ODI to ODI connections of more than one drive may need a cable with 50-pin connectors in between. Also, the last drive in the SCSI or ODI daisy-chain should have the SCSI/ODI termination, while all the others in the daisy-chain should have them removed. For example, if you have a daisy-chain containing DUA0: as master drive and DUA1: as slave drive, the controller board of DUA0: should have the SCSI termination while the ODI termination must be removed from its drive control board. The drive control board of DUA1: should have the ODI termination, and the SCSI controller board and the EDAC board (last two boards) of this drive must be removed.

Following are the jumper settings in the Optimem drives.

With configuration #40:

**a. Controller boards**

Drive DUA0: SCSI ID = 0; no SCSI bus terminator  
Drive DUA4: SCSI ID = 1; SCSI bus terminator

Both the drives should have jumpers 3 (parity enable) and 5 (hardware revectoring enabled) set.

PRECEDING PAGE BLANK NOT FILMED

**b. Drive control boards**

Drives DUA0: and DUA4:	LUN = 0; no	ODI bus termination
Drives DUA1: and DUA5:	LUN = 1; no	ODI bus termination
Drives DUA2: and DUA6:	LUN = 2; no	ODI bus termination
Drives DUA3: and DUA7:	LUN = 3;	ODI bus termination

**Note:**

The SCSI controller board and EDAC board (the last two boards toward the back of the drives) of drives DUA1, DUA2, DUA3, DUA4, DUA5, DUA6, and DUA7 *must* be removed.

With configuration #41:

**a . Controller boards**

Drive DUA0: SCSI ID = 0; SCSI bus terminator

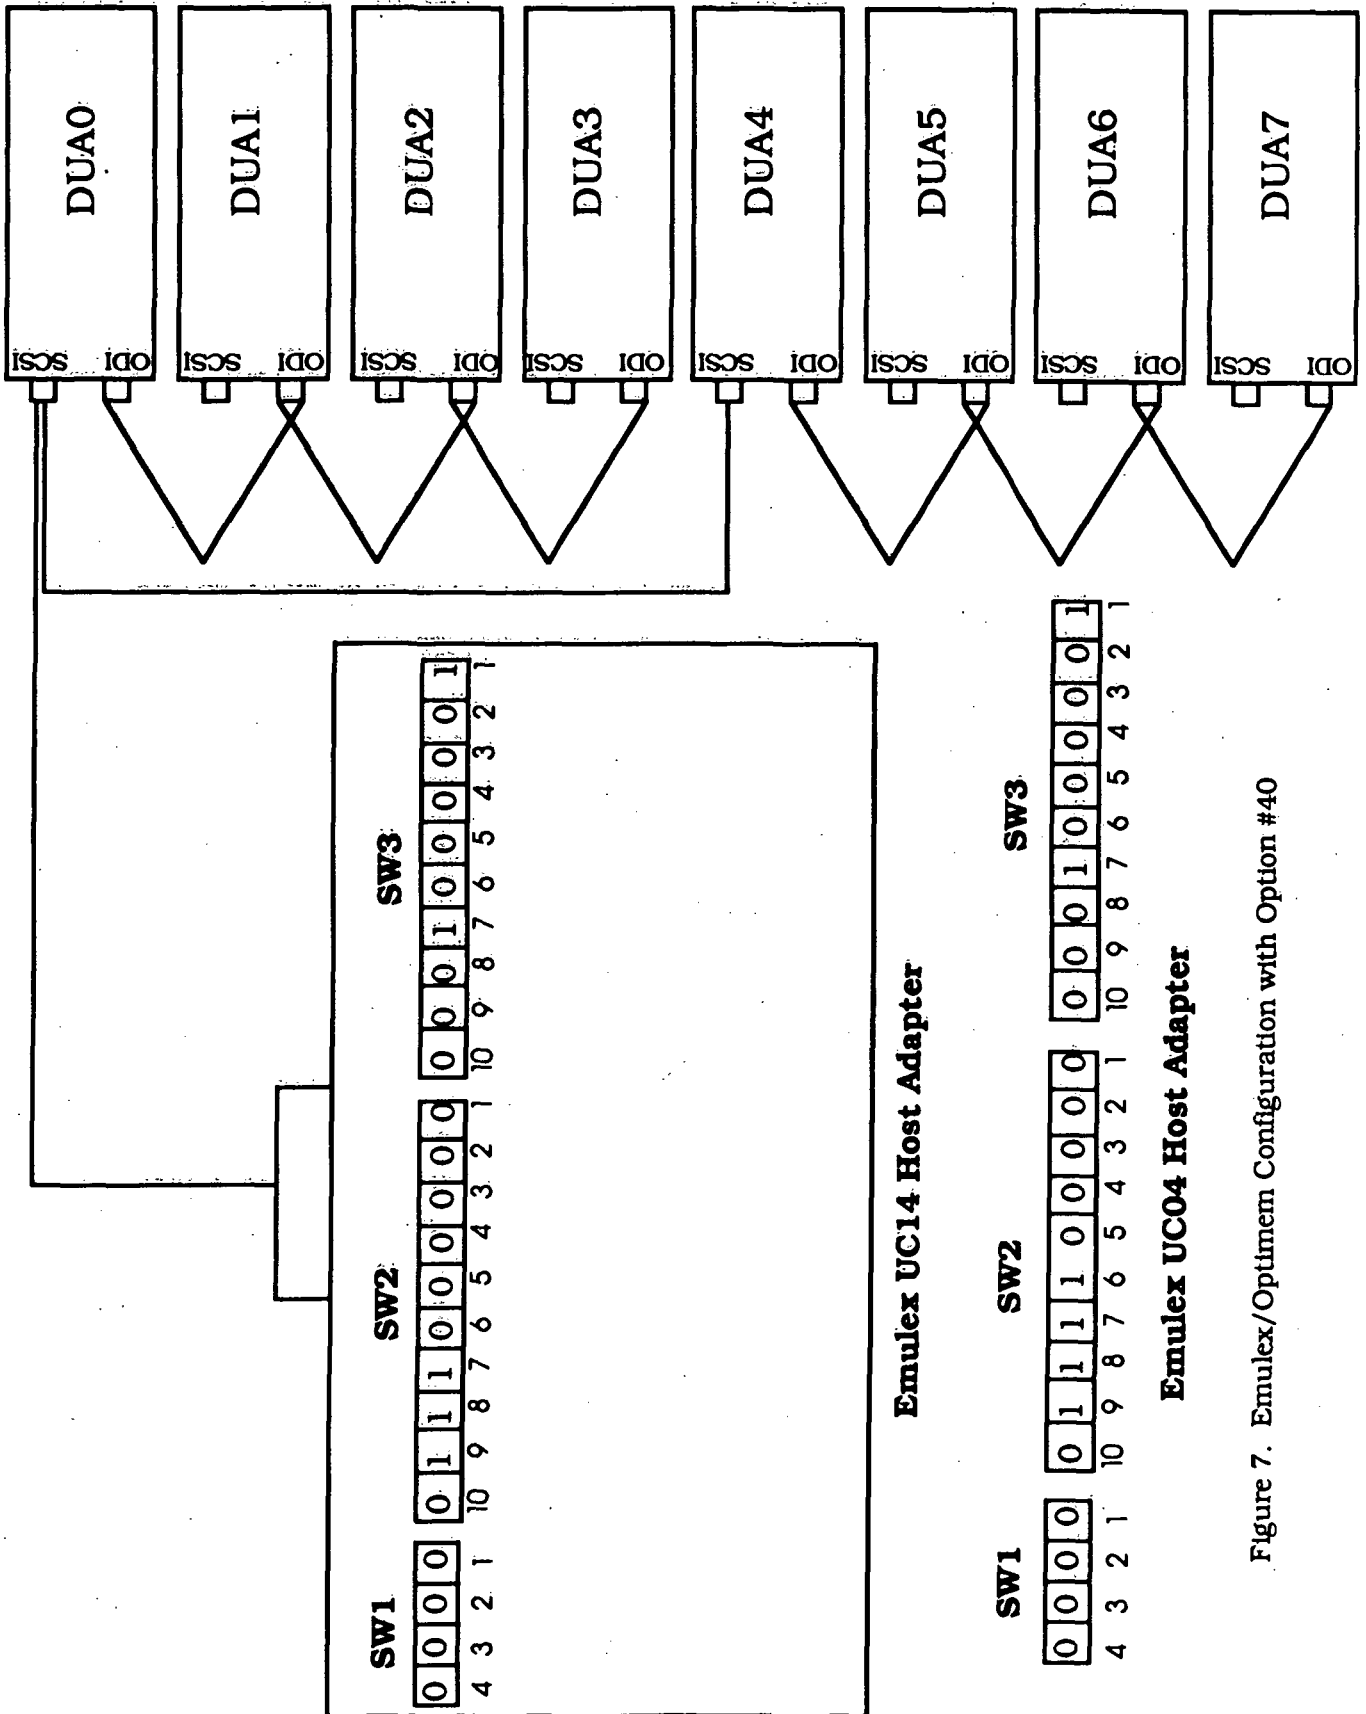
This drive should also have jumpers 3 (parity enable) and 5 (hardware revectoring enabled) set.

**b. Drive control boards**

Drive DUA0:	LUN = 0; no	ODI bus termination
Drive DUA1:	LUN = 1; no	ODI bus termination
Drive DUA2:	LUN = 2; no	ODI bus termination
Drive DUA3:	LUN = 3; no	ODI bus termination
Drive DUA4:	LUN = 4; no	ODI bus termination
Drive DUA5:	LUN = 5; no	ODI bus termination
Drive DUA6:	LUN = 6; no	ODI bus termination
Drive DUA7:	LUN = 7;	ODI bus termination

**Note:**

The SCSI controller board and EDAC board (the last two boards toward the back of the drives) of drives DUA1 through DUA7 *must* be removed.





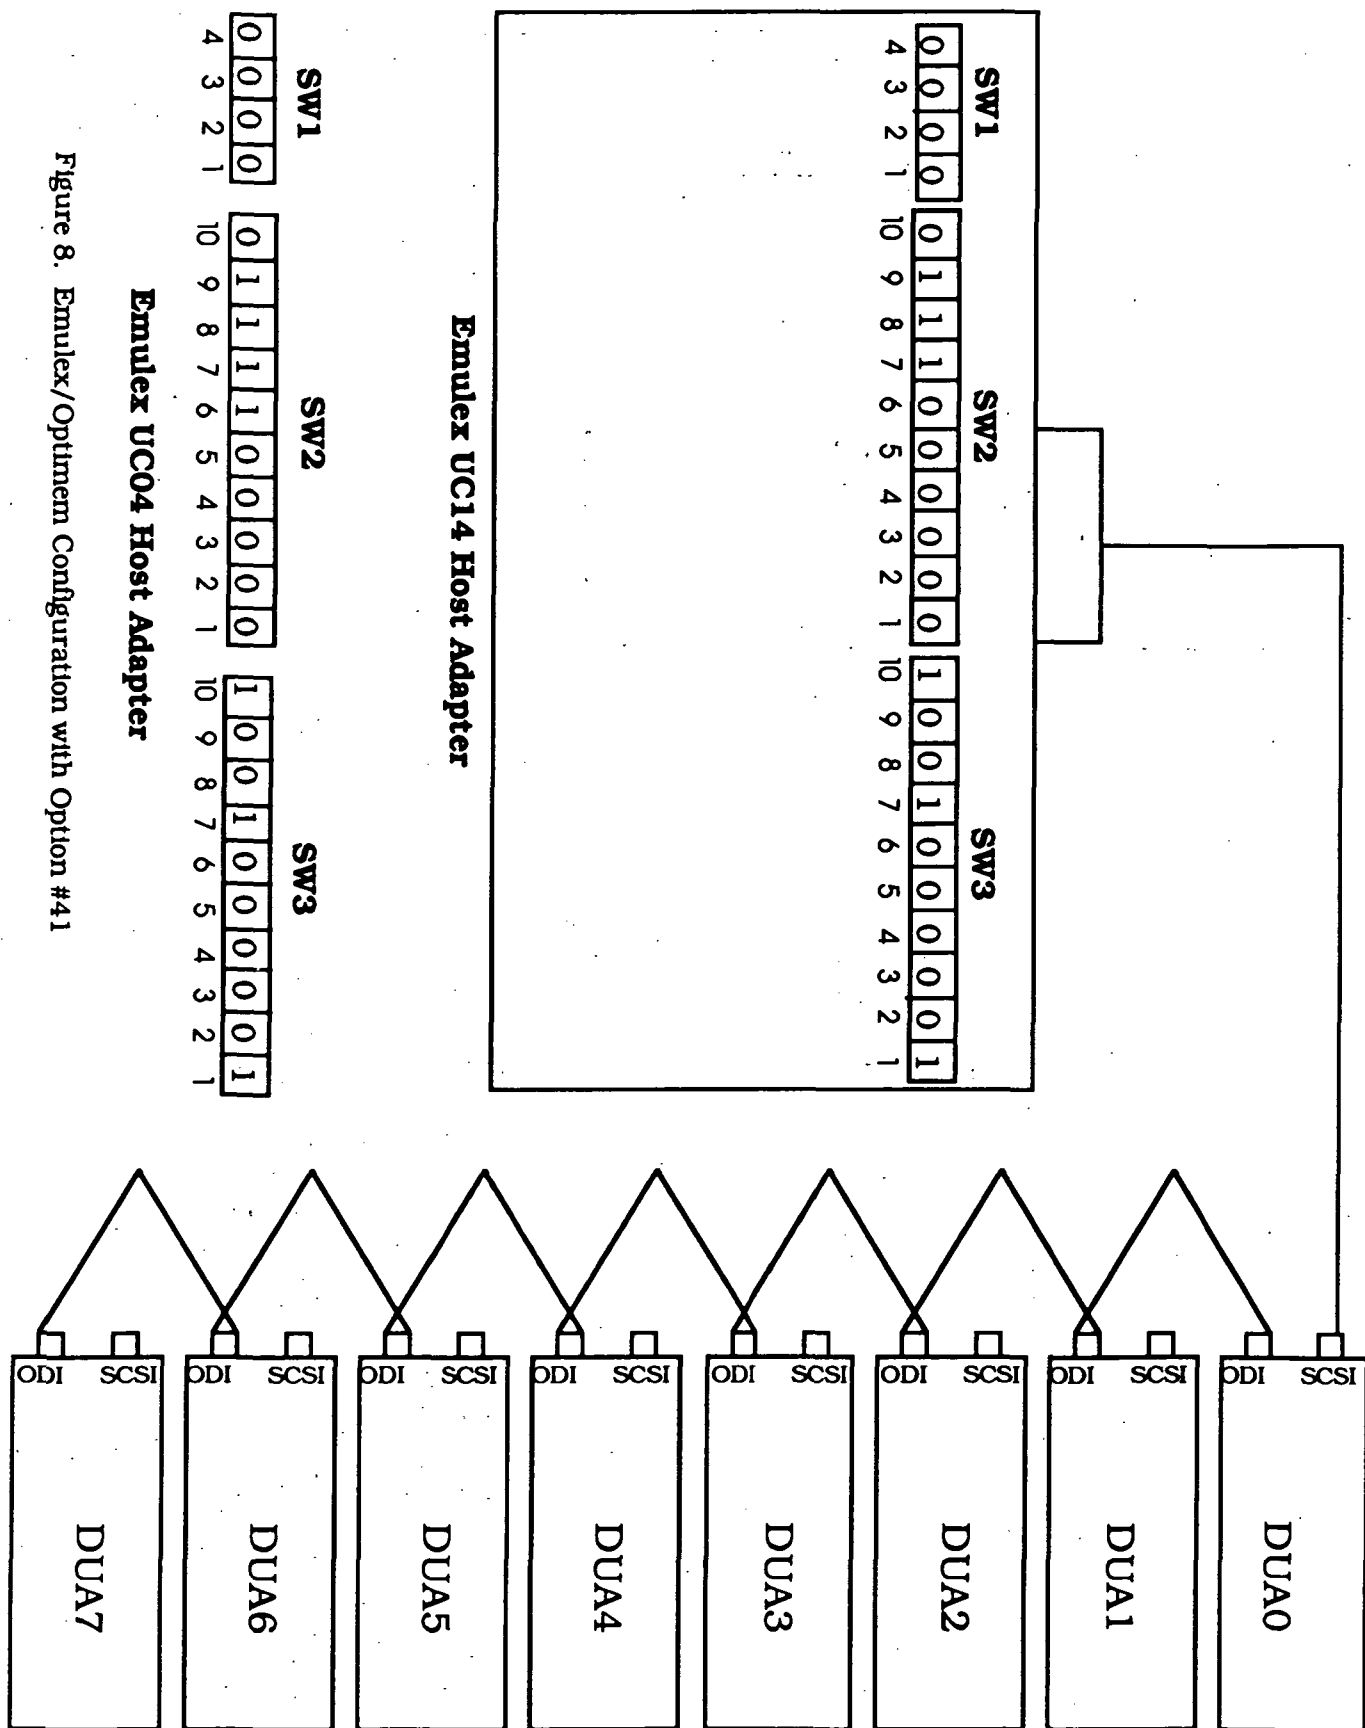


Figure 8. Emulex/Optimem Configuration with Option #41

This section contains a collection of hints to help troubleshoot problems during Emulex controller and Optimem drive installation, RCT initialization, disk mount operation, and SOAR operation.

## **20.1 Installation**

First read manuals (Optimem, Emulex, SOAR).

### **Emulex Controller Installation**

- Check switch settings (Emulex Technical Manual, pages 4-8).
- CSR ADDRESS (Emulex Technical Manual, pages 4-9).
- NPG signal jumper (Emulex Technical Manual, pages 4-13).

### **Optimem Drive Installation**

- Single drive/daisy-chained.
- Verify switch settings.
- Jumper 5 on the Optimem controller board must be set in order to activate the automatic bad block relocation (Optimem Interface Manual, page 6).
- Verify error codes.

## **20.2 RCT Initialization**

- Must be executed standalone.
- RCT initialize as many disks as possible at once.
- If any steps fail, start over from the beginning.

## **20.3 Mounting platters**

- In case of failure (fault 11/12H), try to mount disk a few times.
- If another drive is available, try to mount disk on another drive.
- If a known reliable disk (test disk) is available, try to mount that disk on the drive.
- Check disk for visible dirt.
- If Invalid Media Format error occurs, then RCT initialization is missing or failed.

## **20.4 SOAR**

- Read manual thoroughly.
- The SOAR QSDRIVER requires the VMS DUDRIVER or DSDRIVER.
- Don't use VMS mount to mount an optical disk that has not been OCLOSED, unless the /FOREIGN switch is used.
- Don't use VMS copy to write to an optical disk.
- The magnetic work file requires contiguous space on the magnetic disk.

- Determine the appropriate cluster size for your application. A large cluster size will improve the data transfer rate, but, if many small files are copied to the disk, then a large cluster size will result in wasted disk space.

## 20.5 Additional Precautions

- In case of an electrical storm, turn off the drive.
- Installation of fan(s) in the rack containing the drive doesn't hurt.
- If all the front panel lights are blinking, don't panic! Turn the power off and on again.
- Never issue the standard VMS MOUNT command unless you are **absolutely certain** that the OCLOSED platter you wish to mount is hardware write-protected. Otherwise, you **will** corrupt the optical platter.
- Never use the DELETE command to eliminate files from your SOAR directory. Always use SET FILE/REMOVE. Otherwise, you **will** corrupt blocks and possibly good files on your platter.
- Never use the COPY command to write to your optical disks.
- Never use the INITIALIZE command on your optical disks.
- If you are using DISKEEPER or any other disk defragmentation software, be certain to exclude all of your magfiles from being relocated on your magnetic disks. Failure to do so **will** result in corruption of the optical disk and possible corruption of part of your magnetic disk.